

LOAN DOCUMENT

PHOTOGRAPH THIS SHEET

LEVEL

INVENTORY

New World Vistas...
Information Technology Volume

DOCUMENT IDENTIFICATION

1995

DISTRIBUTION STATEMENT A

Approved for public release;
Distribution Unlimited

DISTRIBUTION STATEMENT

ACCESSION FOR	
NTIS	GRAM
DTIC	TRAC
UNANNOUNCED	
JUSTIFICATION	
BY	
DISTRIBUTION/	
AVAILABILITY CODES	
DISTRIBUTION	AVAILABILITY AND/OR SPECIAL
A-1	

DISTRIBUTION STAMP

19960618035

DATE RECEIVED IN DTIC

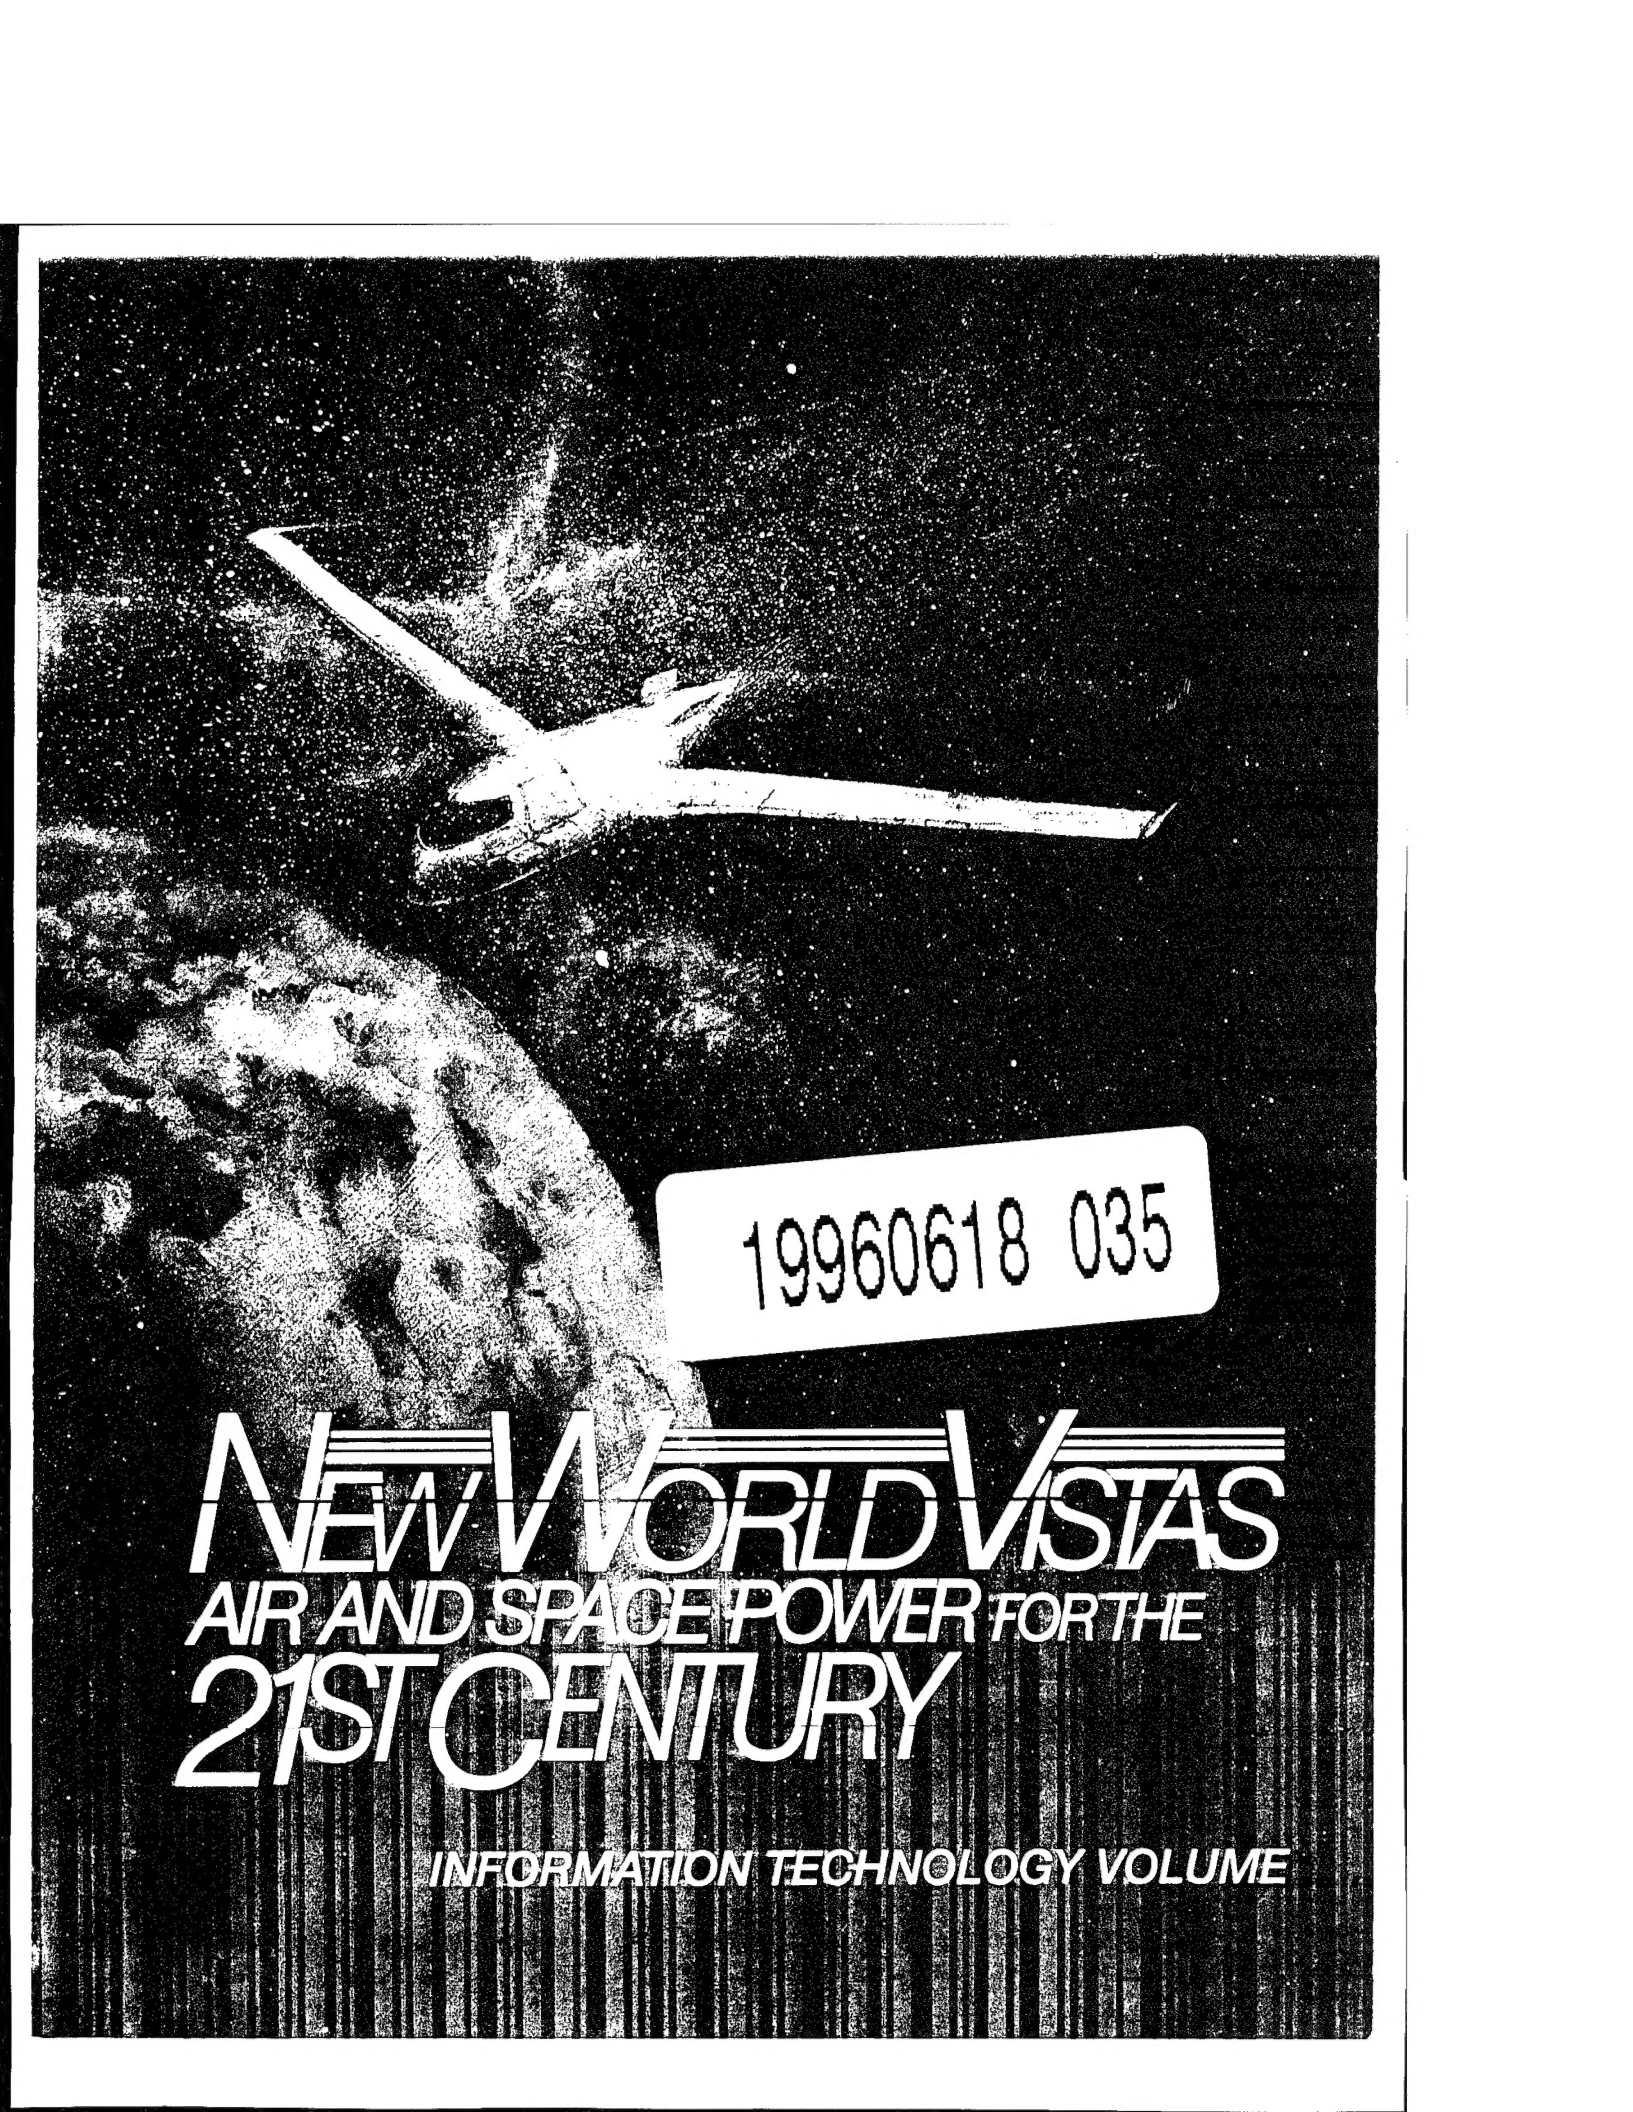
REGISTERED OR CERTIFIED NUMBER

PHOTOGRAPH THIS SHEET AND RETURN TO DTIC-FDAC

H
A
N
D
L
E

W
I
T
H

C
A
R
E



19960618 035

NEW WORLD VISTAS

AIR AND SPACE POWER FOR THE
21ST CENTURY

INFORMATION TECHNOLOGY VOLUME

NEW WORLD VISTAS

**AIR AND SPACE POWER FOR THE
21ST CENTURY**

INFORMATION TECHNOLOGY VOLUME

DTIC QUALITY INSPECTED 1

This report is a forecast of a potential future for the Air Force. This forecast does not necessarily imply future officially sanctioned programs, planning or policy.

Dedication

Major Hubert B. Fisher, Jr.

This information technology report is dedicated to Bob Fisher, a technology warrior for the 21st century. On April 17, 1995, while enroute to a Scientific Advisory Board meeting, a C-21 aircraft carrying Major Fisher crashed. However, this report completes the efforts he placed in progress and hopefully, in his spirit, will maintain our technological preeminence and peace for the 21st century.

Prologue

Consulting the Experts

“It’s hard to predict, especially the future.”
(sometimes attributed to Neils Bohr)

“(Concerning) engines of war, the invention of which has long since reached its limit, and for the improvement of which I see no further hope in the applied arts...”
(Sextus Julius Frontius, Roman engineer writing in ancient times)

“The advancement of the arts, from year to year, taxes our credulity and seems to presage the arrival of that period when human improvement must end.”
(Henry L. Ellsworth, Patent Commissioner, in a 1843 report to Congress)

“Man’s flight through life is sustained by the power of his knowledge”
(engraved in marble at the United States Air Force Academy)

“In the development of air power, one has to look ahead and not backward and figure out what is going to happen, not too much of what has happened.”
(Billy Mitchell: Winged Defense, 1924)

Contents

Dedication	iii
Prologue	v
1.0 Introduction and Executive Summary	1
2.0 Communications.....	16
3.0 Personal Computing	23
4.0 Human-Computer Interaction	31
5.0 Intelligent Software Agents	38
6.0 Information Access technology	45
7.0 Artificial Intelligence	57
8.0 Computer-Aided Planning	66
9.0 Modeling And Simulation	69
10.0 Software Development.....	74
11.0 Software Infrastructures and Standards	82
12.0 Computer Hardware and Architectures	87
13.0 High Assurance Systems	92
14.0 Collaborative Computation.....	96
15.0 Business Applications	103
16.0 IT and Air Force Organization and Education	108
17.0 Epilogue: Principles to Think With	112
Appendix A Information Technology Panel Charter	A -1
Appendix B Panel Members and Affiliations	B -1
Appendix C Forecasting Tables	C -1
Appendix D Evolution and Breakthrough	D -1

Illustrations

Figure 1 Information Technology Dependencies	2
Figure 2 Technology Acceptance - S-Curves	24
Figure 3 Internet Growth	26
Figure 4 SRI's Open Agent Architecture	36
Figure 5 SRI's Open Agent Architecture	41
Figure 6 Transforming Data to Information	46
Figure 7 Interaction of Data and Knowledge	47
Figure 8 Network services from Client-Server to Mediation Layer	48
Figure 9 The Mediation Layer	49
Figure 10 Mediator Design Principle	52
Figure 11 A Mediator has an Owner	53
Figure 12 Integration at Two Levels	54
Figure 13 Allocation Flexibility	54
Figure 14 Some Features of Mediators	55
Figure 15 F-22 Integrated Weapons Systems Database	56

Tables

Table I Motivations for More Natural Human-Computer Interaction	32
Table II Important Types of Human-Computer Interaction	33
Table III Motivations for the Use of Software Agents	39
Table IV Obstacles in the Development and Use of Software Agents	39

1.0 Introduction and Executive Summary

Some military analysts say that warfare at this time is undergoing a revolutionary change. If indeed there is such a revolution, the central element of it is information technology—the technological revolution in computers and communications. It is brought about by an underlying rate of technological progress that is unprecedented in the history of technology and military doctrine.

What makes a technological revolution? Analysis of other technology revolutions tells us that one power of ten of technological change makes a revolution. For example, automobiles transport us ten times faster than walking, and jet planes ten times faster than autos. It took half a century or more for these revolutions to change the face of modern life.

The wealth and security of this nation are increasingly driven by information and are based on knowledge. The engines of knowledge and information, computers and communications, have been doubling their power per dollar every twenty months. The pace is still accelerating. Over the past three decades this compounds to not one or two factors of ten, but six—a millionfold change. And the more powerful these engines become, the smaller they become and the more energy efficient. If this were myth rather than reality, one might say information technologists own a magic wand.

The global village—the wired world of people with shared interests—has been formed because of low cost communications. Computers are so cheap and useful that they are in everything. A tiny but powerful computer sits inside a weapon, guiding it with astonishing precision. Modern warplanes, like the B-2 or the F-22, are as much flying computer networks as they are flying airframes. The remarkable capabilities of modern radar and other imaging devices are as much derived from their computers as from the physics of their sensors.

It is a revolution when computer-assisted logistics planning for a force projection can be done in hours or days rather than weeks. It is a revolution when an accurate situational analysis can be made and given to a warfighter in seconds rather than minutes or hours.

It is a revolution when differences in IT capabilities can contribute in such a major way to a decisive military victory (Desert Storm)—a revolution in which the term “information warfare” is added to the list of future necessary combat capabilities.

The task of the Information Technology panel is to project the visible trends of the continuing revolution in information technology and, where projection fades at the horizon, to envision further progress. We have done this in two ways.

First, systematically we have surveyed the areas of IT work. Examples are communications, computer system architectures, the interface between computers and people, software and the technologies for its development, the emergence of artificial intelligence software that emulates human-like thought processes, software that learns and adapts itself to user needs, technologies for crypto-secrecy and for assured access to systems and networks, and several more. Appendix E contains a sample of IT “futures” of interest to the Air Force.

Second, we have projected and envisioned specific achievements, stretching out over twenty years or more—highlights of the information future. Some are evolutionary, “big wins” with high probability of being achieved. Others represent discontinuities; we do not know if they

will arrive but if they do, their impact will be revolutionary. Still others represent technological, educational and organizational concerns for the future of the Air Force in the era of the information revolution.

Military needs no longer drive this revolution. The good news is that often we can buy off-the-shelf hardware, software, and communications that are much better than, and very much cheaper than, what we can have custom-built for us. The Air Force is challenged to adapt to this new way of doing business, and to benefit from the best that commercial technology can offer (just as our friends and enemies can). But some information technologies that the Air Force needs will not emerge from the commercial marketplace. Our panel made judgments about what these will be as a set of recommendations for continued Air Force and DoD R&D funding priorities for information technology. Our panel also points out where the Air Force can benefit from starting to rethink right now how information technology can improve its weapon system design, acquisition, management, education and career development processes.

What follows are summaries of various IT focus areas, concluding with our recommendations.

Organization and Education					Assurance
Business Functions					
Collaborative Computing					
Personal Computing					
Human Computer Interface					
Software Development	Artificial Intelligence	Agents	Planning	Modeling & Simulation	
Information Access					
Hardware and Architecture					
Software Infrastructure					
Communication					

Figure 1 Information Technology Dependencies

Communications

In the field of information technology it is especially important the Air Force ride the commercial curve. The pace of technology here is arguably faster than elsewhere, and there is an enormous non-military market that keeps prices spiraling downward so long as the needs are not specialized.

The key to information technology is distributed systems, enabled by high speed computer networks. Historically the need for capacity in digital networks has doubled annually,

though the applications that require growth in capacity have not been predictable. Thus it is evident the Air Force must plan for a growth in its networking capacity by about three orders of magnitude in the next ten years, and must reconcile itself to the expectation the uses of that future network cannot be foreseen. Although the prices of commercial communications have also shown an exponential decrease with time, it is not obvious dedicated Air Force networks, e.g., satellite systems, will show this same decrease.

In the near future the work will be networked with very high capacity optical fibers, using wavelength division technology with enormous capacity. The most cost effective communications will rely on fiber networks. However, in many forward locations, the Air Force will not have access to these fiber networks, so the critical problem for the Air Force will be communicating across the approximately 100-mile gap between its forward locations and the world fiber network. The other critical Air Force communications problem is enabling broadband communications into tactical aircraft from ground points not in line-of-sight to those aircraft. The Air Force will need to use satellites, airborne relays, indigenous wireless systems, and deployable wireless systems to cross those critical gaps.

At the current rate of growth most of the world's population will be connected to Internet soon after the turn of the century. The Air Force must plan for the use of this capability to optimize the efficiency of its operations, just as any other business. However, the Air Force must develop specialized skills in the defense and attack of computer networks. It is likely virtually unbreakable encryption technology will be widely deployed commercially and both network attack and defense will be raised to very sophisticated levels. Indeed, skills in the use of information technology will be the differentiating factor in the future for business, and this will likely be true in the military arena as well.

Personal Computing

"Personal Computing" is a category encompassing the information devices likely to succeed today's personal computers as commonplaces in personal and business lives. Thus this is not a category defined by technology per se; rather it reflects the convergence of technology trends from other categories, expressed in commercial products. This category is essentially one of packaging—what novel forms will various components and software be combined into?

In short, this category represents the trajectory of mass market information tools over the forecast period. Understanding this trajectory allows for exploitation of the economies of scale created by commercial deployment—including lowered cost, increased reliability and expanded understandings of systems performance. For example, if Pentagon planners a decade ago knew what Sun, Apple and SGI would be building today, how would that have affected planning? However, such exploitation is not without risk. Incorporation of complex commercially available systems may expose users to unknown risks of penetration and reliability failures.

Finally, it is vital to acknowledge the crucial role played by different age cohorts in defining this sector. The first PCs were invented by individuals who grew up in a mainframe time-sharing world, and thus it is no surprise that the first PCs (e.g., DOS) were near-perfect simulations of what a timesharing experience would be like at a local level. Now, a new generation raised on PCs is about to enter the workforce with ideas of its own. They are certain to redefine the PC sector as profoundly into something utterly new.

Human-Computer Interaction

The next twenty years will see two important ways computers will come to serve us. The first and undoubtedly the most compelling vision is their migration from a hands-intensive tool to a delegatable assistant. The ever-affordable surplus of computing cycles will be more and more dedicated to making the computer not just user-friendly but user-like.

The more straightforward use of computers today is as a framework within which application programs are run. The processors, the connectivity, the operating systems are becoming invisible, subordinated to the applications programs at hand. Within a decade those, too, will begin to sink below the conscious surface, yielding to task-oriented computing as the next level of work abstraction. To make this possible, the way people interact with computers will take on many of the attributes used when they interact with other people. Over the next few years, speech and handwriting recognition and understanding software will permit a wide range of computer interaction now confined to the keyboard and mouse. That does not predict their demise, however.

The second compelling use of computers will be to augment the physical, not just the intellectual human. Computers offer a very flexible, yet precise means to transform what we are equipped naturally to use into realms of functionality we would have never dreamed beforehand. While the range of such use may not be as broad as more mind-centered work, there are many situations where our dimensionality, our sensory package, our reaction times, our motor sensitivities, or simply our remoteness are inappropriate to the task at hand. The most general concept of such transformation is telepresence, providing human interaction potential to otherwise inaccessible environments. Emerging now are computer-mediated tool systems that permit exquisite sensory and tactile presence from a remote location and with dimensionality scaling where needed. The handling of dangerous material, the teleoperation of large- and small-scale systems, and remote surgery are a few examples.

At the intersection of the mind- and physical-centered worlds of computer use lie new fields like augmented and virtual reality. Augmented reality is the overlay and supplementation of real world, real-time events with simulations intended to instruct or focus the attention of the participant. Virtual reality involves a totally synthesized environment for the participant in which, ideally, all senses are used. Virtual reality systems are, of course, with us now but affordable systems that closely match our sensory capabilities are still a decade or so away.

Intelligent Software Agents

Though the notion of software-based agents have existed since the early 1980s, they are just now beginning to emerge as functional parts of software systems. The lack of a precise definition of agents, unfortunately, gives the developer and marketer wide latitude in just how much or how little functionality is present. At their best, software agents are capable of representing a user or owner in the accomplishment of specified tasks without his/her having to prescribe how it is to be done. Retrieving specific information is the most frequently suggested role for agents, but the complete range of functions is extremely broad.

We define agents broadly, as program that: have an owner, have some degree of autonomy, be goal and not means driven, and be able to create other processes or data. Given the predicted state of development over the forecast period, any use of agents in safety-critical situations must

be guarded. The intelligence displayed by agent programs will evolve almost indefinitely. For the moment we define three levels: (1) directed-action agents have fixed goals and limited ability to deal with the environment and data they encounter; (2) reasoned-action agents have fixed goals and an ability to sense both environment and data and take reasoned action; and (3) learned-action agents accept high-level tasking and are capable of anticipating user needs based on general guidelines and can issue themselves new goals in the process.

Over the next decade agents will be able to provide such functions as advising and personal assistance. Advisory agents will be able to monitor a situation and give feedback and recommendations, such as in instructional settings. Personal assistants will most likely appear as managers of human-computer interaction, offering user assistance in specifiable tasks such as electronic mail, calendars, conference set-up, information or more general searches. They will come to represent the user in a wide variety of transactions. "Traveling" agents will be aware of the vast, network-based world of information and its indexing. Under specific requests and general constraints, they will retrieve information of all kinds, potentially operating within remote servers or hosts. They will be needed because the inventory of sources is so large, its quality so uneven, and its costs so varied that humans will not easily abide such searching.

Information Access Technology

Information Access Technology (IAT) provides the framework and foundation for accessing and transporting relevant information to the Air Force decision maker. Central in this task are the conversion of data from the many heterogeneous data resources available into an integrated information base that provides situation awareness. Additionally, IAT needs access to software for symbolic fusion and summarization. Advanced IAT services must also have access to simulation software, so the decision maker can develop and compare multiple courses-of-action within the system in terms of their potential results, resource consumption, and risks.

A critical task of IAT services is the reduction of data to a manageable amount of essential information. Integrating many streams of data further increases the volume of data delivered. It is crucial to avoid information overload to the decision-maker (whether an airman on the maintenance line, a pilot on a mission, or the commander of a joint task force), by presenting only relevant information at the appropriate level of abstraction. Information overload also has to be avoided when lower bandwidth links are to be used, often in the last 100 miles to the warfighter.

Two types of IAT modules are usefully distinguished: mediators and facilitators. A mediator performs such tasks as summarization, abstraction, and fusion at the behest of its owner. Mediating modules can either provide information according to a pre-planned schedule, or in response to a request. A facilitator, by contrast, is a module that searches the network for sources of data, and delivers newly discovered data to a mediator and its maintainer, or directly to an end-user recipient.

Multi-level, distributed IAT systems are enabled by the rapid progress in communication hardware and infrastructure technology. To gain the maximum benefit from such systems, we need experience with IAT concepts involving modern technology so rapid assembly of IAT systems will be possible, e.g., to assemble mission nets for new situation and rapidly assembled task forces.

Artificial Intelligence

The information technology called artificial intelligence (AI) is used to build software that performs the kinds of information processing tasks humans perform so well—using knowledge to solve problems, analyze situations and decisions, recognize patterns, form hypotheses, learn new knowledge, understand human speech and language, and so on. To date the technology has demonstrated limited but significant real-world uses in both commercial and military environments and is poised for widespread and powerful uses in the next two decades.

The key limitation to the AI technology has been a limitation of what AI programs know. To perform well, AI software must be more than narrowly smart, it must be broadly educated. The narrow specialty knowledge bases of today will be expanded in the next two decades to huge knowledge bases with tens of millions (later hundreds of millions) of pieces of knowledge and know-how. Much of human knowledge will be encoded for use by both people and AI software, and will be held and distributed on the Internet. The key bottleneck in doing this—acquiring the knowledge for encoding—will be overcome by technologies now emerging that allow the automation (or semi-automation) of the learning process.

We expect that accurate understanding of human language in text and speech will reach the 95% level in ten years, for areas of limited scope and will handle broad areas of human discourse within twenty years.

Early experiments suggested that AI technology can provide a powerful expanded framework for the problem of information fusion for situational awareness, allowing fusion not just of multiple streams of sensor data, but of situational knowledge, experience, and context as well. This fusion of signals with symbols will be a key to solving information overload, as well as understanding image data and intelligence collections.

Finally, AI application will bring big changes in affordability and rapid response in the areas of designing, simulating, building, and maintaining equipment and software.

Computer-Aided Planning

Improvements in planning and scheduling over the next twenty years will provide revolutionary advances in the speed, efficiency and effectiveness of Air Force operations. Existing techniques, such as constraint-based planning, when widely applied, will provide orders of magnitude improvements in operations planning. Additional advanced capabilities will arise from capturing plan rationales (i.e., the reasons why actions are being taken), enabling for example the speedy modification and re-use of existing off-the-shelf plans. Advances in decision theory and reasoning about uncertainty will enable execution-time plan modification even in the highly uncertain environments characteristic of conflict. Finally, a variety of increasingly powerful multi-agent planners will permit the integration of concurrent multiple viewpoints and, eventually, permit planning to be done continuously by multiple teams of people and software agents working together.

Modeling and Simulation

Computer modeling and simulation provides exploration of alternatives for the full spectrum of Air Force activities, from research and development through analysis, acquisition, test, evaluation, production and logistics to education, training, and operations.

Computer modeling and simulation has been evolving its broad range of Air Force capabilities since the first use of computers for ballistics applications in the 1940's. As a result, the current Air Force inventory of independently-developed computer models and simulations has formidable interoperability and compatibility problems. These have been overcome in some limited domains (e.g., flight dynamics models) and in some medium-scale training-oriented distributed interactive simulations such as Simnet. They have also been overcome in point-solution demonstrations and exercises linking wide varieties of models, simulations, real equipment, and operators. However, significant broad, regular-use operational issues such as interoperability, rapid configuration, and verification, validation, and accreditation (VV&A) are just beginning to be addressed.

By 2005, with a significant level of Air Force effort, basic large-scale interoperability support will be available, including consistent data definitions and fixed interaction protocols. By 2015, this will extend to dynamic interaction and interoperability agents. By 2005, simple user languages enabling rapid composition of models and simulations will be available. By 2015, these will extend to support automatic configuration of models, and simulations that address decisions such as choice of airlift capabilities.

By 2000, the Defense Simulation Internet will provide broadband support of point-solution "Louisiana Maneuvers" scale simulation. By 2010, it will be operationally robust and able to support regular exercises at this scale. Currently, VV&A technology consists of basic test suites and simple assertion checking (e.g., of conservation of energy, resources, etc.). By 2005, with Air Force investment, this can expand to simple mission-domain model checking and built-in-tests. By 2015, this VV&A technology can expand to domain model checking using automated agents and dynamic built-in-test, achieving much higher levels of credibility.

By 2015, the resulting modeling and simulation capabilities will enable combat operations options to be credibly simulated before and during combat, greatly increasing combat effectiveness. The same capabilities will enable continuous two-sided exercise of information warfare capabilities, honing Air Force pre-eminence in this critical area. With appropriate attention to acquisition restructuring (e.g., virtual competition groundrules), the technology will enable virtual system acquisition, or flexible migration from virtual to actual combat systems, with complementary closed-loop combat system exercise and improvement across the system's life cycle.

Software Development

Software development technology provides the Air Force with its best defense against tomorrow.

The other technologies addressed in New World Vistas are primarily focused on ensuring that on any given day in the future, say July 4, 2025, the Air Force is better prepared to prevail in combat than its adversaries.

However, suppose that on July 5, 2025, the Air Force detects that an enemy has discovered and is ready to exploit a critical weak link in Air Force combat capability. What technology can best enable the Air Force to repair the weak link and disseminate the fix to its full complement of forces?

An electronically disseminated combat system fix would best fit this need. The most powerful and flexible electronic fixes by far are changes to the combat system's software. Enabling the Air Force to design, develop, test, and deploy these fixes is the province of software development technology.

Software development technology consists of methods, processes, tools, and assets enabling faster, cheaper, and better development of computer software. Methods, such as object-oriented design, configuration management, and Cleanroom methods, enable improvements in software design, verification, version control, and quality assurance. Processes, such as incremental, evolutionary, spiral, design-to-cost/schedule, and product line management processes, enable more efficient orchestration of the methods. Tools, such as design, code generation, test, product and process management tools, can completely or partially automate software development functions. Reusable assets, such as specifications, plans, components, and test cases, enable more of a software system to be composed of already-developed artifacts.

As with other information technologies, software development technology is increasingly paced by commercial investments. From an information warfare standpoint, this must be considered a technology leveler with respect to the Air Force and its adversaries. Not using commercial software development technology would be a competitive disadvantage for the Air Force, but so would a purely reactive use of commercial technology.

Software Infrastructure and Standards

The breadth of computer use in commerce permits the military to exploit the massive investment in software being made for commerce, manufacturing, education, and entertainment. This investment will continue to drive progress, the benefits of which will be available to all: commerce, the military, our friends, and our adversaries. In this environment, speed is the key to remaining current: the Air Force must ensure that its acquisition processes enable it to move quickly in acquiring and exploiting new software.

The Air Force has a need for standards that enable interoperability that far outstrips what is likely to be produced by the commercial marketplace where individual vendors have interests in developing proprietary systems.

Progress Due to Commercial Interests.

Commercial interests are causing a rapid convergence to a smaller number of distinct software (SW) systems. We can project the ongoing commercial investment will be limited to one operating system for mainframe computation, a family of operating systems clustered around UNIX, and perhaps two or three families of systems on personal computers. There will also be broad use of standards among those systems, for instance SQL to enable UNIX systems to access mainframe databases, and object-oriented protocols to enable personal computers to access object-structured resources kept on UNIX-based servers.

These interoperability standards will continue to grow in breadth and capabilities. Similarly, features made available in one system will be rapidly replicated and superseded by features provided by a competitor. Even though the pace of change may slow down somewhat, there is little doubt throughout the next 25 years the commercially available software

infrastructure will provide the foundation for most DoD and Air Force systems. In mission-critical applications, because of the need for security and trustworthiness, DoD will lag behind the curve with respect to features and cost-performance ratios. This lag should be minimized, so DoD SW availability and its capability to exploit modern hardware does not lag excessively behind adversaries, who may well be willing to proceed with fewer constraints on security and trustworthiness.

Capabilities Appropriate for Government Support Funding.

Users need to retain the flexibility to change platforms and use systems composed of hardware and software obtained from multiple vendors, including competitors. Standards conformance should be supported by government, to assure fair and early dissemination of sharable technology.

While NIST has the primary role for standards, there is a role for DoD research to support standards activities that can accelerate the adoption and availability of sharable technology of interest to the military from commercial vendors.

Conforming military subsystems, integrated into the commercially available infrastructure, will provide the basis for future military information and control systems. They will displace the bottom layer of the stovepipes that characterize DoD system architectures and provide the infrastructure for interoperation of general software development and the global information systems.

Capabilities Requiring Military Investment.

On top of the commercial and national infrastructure will be requirements that require specific military investment if US SW technology is to be maximally effective and superior to the technology of adversaries. Specific capabilities will include:

1. Highly robust real-time SW modules for data fusion and transmission.
2. Integration of simulation software into military information systems.

Computer Hardware and Architectures

The microprocessors made by a few merchant semiconductor manufacturers will dominate computer hardware designs. Since tens to hundreds of millions of these will be manufactured each year, the low cost of 'printing' so many computers will be irresistible to computer makers.

The microprocessors will be arrayed in networks of various sizes to satisfy product specs: uniprocessors, small multiprocessors, super-scale arrays for major computing power, networks of workstations, and so on. Also scaled to the product design will be networking capability linking the common microprocessors.

Except for stand-alone computers in embedded applications, the slogan "the network is the computer" will become a reality. Computer systems architecture will be largely "plug and play." This implies major changes for the system software, something discussed in another section.

Ultra-large-scale integration of circuits on silicon will follow its historical trend of doubling power per dollar about every 20 months for the next decade. At line widths of one-tenth micron, this remarkable path may end, and other solutions to obtain microprocessor power may have to be found. The solutions may come from new materials (e.g., “bio-organic”), new physics (e.g., atomic-level switching), or new designs (e.g., new approaches to parallelism) or from an entirely new concept not yet known (unlikely).

The wide use of just a few families of microprocessors will lead to the kind of defacto standardization that will make commodity items of all peripheral devices. Huge amounts of on-line secondary storage will be achieved by evolving into parallel-organized databases connected to high bandwidth bitways.

High Assurance Systems

Existing information intensive systems are currently blatantly vulnerable: recent studies have shown the domestic electric power grid, financial systems, and telecommunications infrastructure to have between modest and virtually non-existent protection against information-based attacks. Yet basic techniques exist capable of deterring many of these attacks, and continued development and dissemination of known cryptologic technology will be able to provide very high levels of security to individual systems. Attention should be paid to widespread integration of such technology into Air Force software at all levels (networks, operating systems, and applications). Important attention should also be paid to Air Force *policy* regarding cryptology: in particular, we recommend the Air Force employ a key escrow (or similar) system, in order to ensure that internal use of cryptographic techniques cannot provide an impenetrable wall of privacy to unauthorized action by Air Force personnel.

Difficult problems arise in providing high assurance and survivability for larger-scale distributed systems. A significant body of foundational work exists in fault tolerance for modest scales of distribution, but significant work needs to be done to provide assurance and survivability when potentially every computer in the Air Force will be interconnected. Of particular note will be development of techniques for graceful degradation—the ability of a system to provide appropriately selected partial functionality in the face of unanticipated failures, rather than the all-or-nothing functionality provided by today’s approaches to fault tolerance.

Finally, we note with considerable concern while a great deal of attention has been given to the notion of information warfare, there is as yet no notion of rules of engagement in this field. We cannot tell when we are under attack, first because information attacks can be considerably more subtle than physical attacks, but more importantly because we have not yet established any notion of what constitutes a hostile act in cyberspace.

Collaborative Computing

“Groupware” is the label for the notion that computers and information can be used to support business teams, rather than mere individuals (the dominant personal computing paradigm in the last decade). Groupware thus encompasses both the creation of entirely new systems, as well as the pressing into group service systems designed for individual use.

The term “Groupware” was coined in the early 1980s by Peter and Trudy Johnson-Lenz, but groupware systems are yet older. One of the oldest, the Augment system built by Douglas

Engelbart at SRI in the 1960s, also happens to be among the most ambitious, incorporating features such as two-way video, group writing, and shared multiple-window screens.

Groupware products are proliferating today, yet the term remains more a concept than a product category in its own right. In the long run "groupware" as a discrete category will disappear entirely, as group-oriented features are incorporated into virtually every product offered in the PC sector.

Groupware as a product category is growing solidly but modestly. The single largest groupware product is Lotus Notes, which is enjoying considerable popularity among corporations, but still has an installed base which is minuscule compared to the total installed base of networked PCs in corporations. Notes itself does not have a bright future, and the features that make it popular will find their way into more generalized environments such as operating systems, advanced data repositories and even Internet browsers and servers.

Groupware will have its largest impact on organizational structures. Even the most casual glance at business history makes it clear that each time a new information infrastructure becomes available (e.g., railroad, telegraph, telephone) the entities which are ultimately most successful are also the first to reshape their structures in order to gain maximum advantage of the new information conduits. The new networks emerging today are "geodesic" (a term first noted by Peter Huber in the mid-80s in the context of telephone deregulation), that is, global, non-hierarchical and without any central node. It is a safe bet that our organizations will follow suit.

Business Applications

Information technology will change the Air Force way of doing business, thoroughly permeating the Air Force of 2025. From the desktop to the flight line, the business needs of the Air Force will be supported by information appliances.

The business side of Air Force operations will be contracted, not done in house. Outsourcing functions like payroll, personnel and property management will help meet down-sizing and budget constraints.

Information technology will provide the means to commercialize and to more effectively accomplish its remaining business functions. Bandwidth on demand for any application will allow telecommuting from anywhere to anywhere. Coupled with extensive desktop computing power, collaborative planning and telepresence will be the norm. This, EC/EDI, and improvements in our modeling and simulation capabilities will dramatically alter the Air Force procurement and acquisition process. Systems will be created using advanced modeling and simulation and electronic engineering and then they will be "test flown" in a Star Trek like holodeck before the Air Force buys or builds the first one.

The contract business of the Air Force will join the rest of the corporate world to become another paperless enterprise. Electronic tethers, in the form of personal communication devices, that use voice, video or data communications and intelligent assistants will keep blue-suitors in constant communication with their offices. The bottom-line, the business side of the Air Force will continue to be no different from the business side of the corporate world and COTS software will meet AF needs.

Organization and Education

The revolution in information and IT will profoundly affect the organization and education of the Air Force. As the industrial revolution created entirely new organizations of mass production, the information revolution will lead to new ideas of what organization and education mean. Further, as air power once significantly changed army and navy doctrine and organization, IT power will change the AF even if the primary responsibility for the Information War mission is consigned to another agency. Organization and education will be vital differentiators for the Air Force in a world of commodity technology.

But what are the characteristics of effective organization and education 25 years from now? The beginnings of an answer can be found in organizations that are today riding the IT bow wave. These enterprises are flatter and organized to produce hundreds of niche targeted products and services. To effectively adapt the Air Force must learn how to strike and defend with hundreds of varied lethal and non lethal weapons from physical ordinance to abstract bits and bytes. A second related characteristic of successful organizations is their attention to front line employees. These workers are empowered with intellectual and informational skills, and the power to act on what they see and observe. Currently, the Air Force is decentralizing the power to act on decisions, but does little *intellectual* and *informational* skill education necessary for sophisticated learning. Instead it tends to rely on training that emphasizes standardization, top down direction, on-again off-again phasing, and physical skill training. These training courses are still useful, but would only be one aspect of a multi-faceted educational system.

New educational processes within this flatter, networked, multi-faceted AF will be more on-going, continuous, and student-centered. Education, defined broadly to include the intellectual and information skills, then will become a vital daily transaction for the organization, in addition to sortie generation. Perhaps contracted mentor-nets involving hundreds of the best educators and thinkers in industry and academe will provide the opportunity for AF “cyber pilots” and “bit jocks” to hone skills, share insights, notice threats, exploit adversary weaknesses, and differentiate our capabilities from hostile organizations.

Organization and education are not distinct, but blended facets of information power are a prerequisite of successful organization. Reinforcing each other they help create a combined system capable of fighting and defending in hundreds of niches. This is similar to how smart terminals within simple networks have replaced large central controlling computer networks.

Recommendations

Prologue

Future development in information technology will come predominantly from the commercial sector. Military needs are a small part of rapidly growing commercial markets. The key drivers are commercial economies of scale in the production of chips, software packages and fiber bitways and from the importance of standards, de facto or official, that are critical to all information transfer.

However, not everything that the Air Force will need in the future will be available in the commercial marketplace. For those needs that we list below, the Air Force must make long term R&D investments. In some cases they represent added functionality and in others they will be

strictly military-unique. These areas will help provide the differentiators so critical to a military future in which both sides have wide access to the same commercial information technology. This is our “focus” message for Air Force IT.

We also have a “defocus” message. We list below areas in which commercial firms will produce products that will be satisfactory for AF needs.

Needed Air Force Investments in long-range R&D (or shared in vestments with other DoD entities or private companies).

1. Information transfer over the backbone-to-mobile-platform link “the last N miles”. Consider intelligent compression, differential updates, high bandwidth directional antennas to create a highly available, minimum bandwidth core.
2. Information fusion comprising both signals and symbolic knowledge.
3. A widely available knowledge web of tens to hundreds of millions of pieces of knowledge. The AF shares in the research costs and develops AF specific knowledge for AF needs.
4. Software architectures that work with AF-specific knowledge, reusable components, safety-critical components, real-time systems, and other military-oriented capabilities.
5. The automatic indexing of images by their semantic meaning in terms of military objectives (intelligent image retrieval).
6. Automatic capture of the rationale for plans during planning activities.
7. Reasoned-action and learned-action agents whose goals are AF-specific.
8. COTS software components and CASE tools enhanced to meet AF combat needs (such as security, survivability, real-time performance, and scalability).
9. Architecture for “just-in-time” information systems and networks “when you need it.”
10. Multi-agent planning software.
11. Software that is survivable, displays graceful degradation.
12. Modeling and simulation for training; system design through acquisition; planning and decision making. This will include component interoperation; believable semi-autonomous forces; transition from virtual to actual system acquisition (in a non-proprietary way); validation verification, and accreditation.
13. Augmented reality—the overlaying of synthetic, spatially synchronized, cues and structures on real-time, real-world activity for training and for the maintenance of complex systems.
14. Human-computer interaction capable of sensory-matched control of all UAVs.
15. The technologies and rules of engagement for information warfare.

-
16. Telepresence, the real-time trans-lation of the human senses and physical dexterity into otherwise inaccessible spaces with possibly “non-human” scaling factors (e.g., very large, very small).

Defocus AF investments from these areas that will be well-handled by the commercial sector:

- A. High capacity communications “backbones”; global telephone networks; world-wide wireless infrastructure, Internet, ATM.
- B. Cryptography routinely embedded in systems
- C. Compression (except intelligent compression)
- D. Multimedia technologies.
- E. Natural Language Understanding, including Speech Understanding.
- F. Computer displays
- G. Data mediators, request facilitators, information broker software
- H. Basic directed-action software agents
- I. Software for the “business” functions of the AF: logistics, personnel, finance, etc.

Next Steps:

1. Air Force laboratories must rethink their information technology R&D programs in the light of this New World Vistas report
2. The Air Force should rethink its advanced weapons system design from the info-centric point of view
3. The Air Force should rethink the education, career path, and reward structure for its officers and airmen in light of the IT future projections in this report
4. The Air Force should rethink its acquisition strategy in the light of advanced IT capability

A Sample of Forecasts in Information Technology Important to the Air Force:

1. All cryptographic codes will be unbreakable, but systems may still be breakable.
2. All computer networks and platforms will be scaled from the commodity components of personal computer and workstations.
3. Computers will understand natural language.
4. Bits-per-second in a backbone will be effectively infinite and of near zero cost.
5. Information transfer over the backbone-to-mobile-platform link will remain bandwidth-constrained.
6. End-user programming can harness large bodies of accepted code.

7. Information fusion will comprise both signals and symbolic knowledge.
8. Intelligent programs will carry out "what-to-do" commands instead of "how-to-do-it" commands.
9. Advances in modeling and simulation can enable significant advances in system design, planning, acquisition, training, execution, and reuse.
10. Human-computer interaction will enable sensory-matched control of all UAVs.
11. A profound shift toward information technology must occur in AF education centers and R&D laboratories.

2.0 Communications

2.1 Information Technology—An Overarching Concern

With widespread proliferation of advanced military weaponry and capability, what differentiates one nation's force from another's may be their relative usage of modern information technology. Clearly the USAF will have to deploy advanced global communications networks in order to have the kind of sophisticated information environment that will be commonplace in the near future. However, at first glance it is not evident that the networking that the USAF will be able to project into the field will measure up to the broadband capability that may be accessible to the average user in any developed nation in the world.

The information industry is moving massively and quickly towards architectures that depend intimately upon broadband communications. Two current trends that fuel this thrust are the client/server model of distributed computing and the move towards increasing use of multimedia information. Shrink-wrapped software will be available everywhere that depends upon high speed communications connectivity. This is the environment with which Air Force personnel will become familiar when working in their homes and in their permanent bases. The question is whether they will have to give up this bandwidth-rich environment when they are deployed to forward positions. If so, then the USAF will be isolated from the natural technology of the information industry, and this has strong implications for strategy, tactics, and training in the future.

Information Technology will be the differentiator in business and in the military. The Air Force must ride the commercial wave. The pace of information technology is faster than the typical acquisition cycle. Open interfaces and standards will dominate. Information technology will be ubiquitous in the world; entities will gain the advantage through skilled use of the technology, rather than through development of unique technology.

2.2 Projecting the Future in Communications

The future in communications arises from the interplay of a number of factors with differing time constants:

- Exponential technology evolutions (e.g., Moore's Law)
- The evolution of standards (e.g., ISDN)
- Economic flywheels (e.g., investment in infrastructure)
- Market chaos and fashion

Some technology is predictable. Moore's law predicts the exponential decline of transistor price and the exponential increase of computer power. There is a similar law for the capacity of optical fiber systems. Compression technology gives steady advances in bandwidth efficiency.

Advances in standardized infrastructures, such as ISDN and ATM, are also predictable. They happen over many years, as the technology must be developed and popularized. We could predict now that ISDN will rapidly disseminate over the next several years as an improved access method for Internet and other digital networks. ATM will take five to ten years to become a ubiquitous platform for high speed multimedia communication. Right now we foresee ATM

being the technology of choice for LANs, for wireless access, and for backbone carriers. However, it appears that a future version of TCP/IP (version 6, or its progeny) will be the main protocol for packet networks. There will be many conversions between IP and ATM.

Another factor in how the future evolves is the economic flywheel. For example, it will require tens of billions of dollars to wire the US for broadband access. This will take perhaps a decade to accomplish.

Finally, market forces, serendipity, invention, and fashion form a fertile chaos from which many futures could coalesce. This is so important that it cannot be overstated. Simply put:

We do not know the future uses of networks—they are inherently unpredictable. The growth of the Internet is an example of this unpredictability. The exponential growth is the envelope of a number of rising new applications, the first of which was electronic mail and the most recent of which is the web. None was foreseen, yet all together they continued a predictable, exponential growth.

2.3 Laws of Networks

The law of network traffic—it doubles annually.

The law of network value—the value of a network grows with each additional user

It is well known that data traffic on computer networks has been growing exponentially. Reliable statistics are kept on the traffic on the Internet backbone that show a steady 100% annual growth rate over the last decade. Since this is the largest computer network in the world, such statistics are rather compelling in thinking about how to project the natural growth of data traffic in Air Force networks. However, there is a complication—the Internet also shows a 100% growth rate in number of host computers. (The growth of actual users is not known.) Thus it is possible that the growth in traffic in the Internet is entirely due to the continual expansion in the number of connected computers. In Air Force networks we would not expect such a growth of computers.

Statistics on the growth of data traffic in fixed-user-population networks are much harder to come by. However, one comprehensive study at the Lawrence Berkeley Labs shows that in their own network they experienced a 100% annual growth rate over a period of years when their size, in terms of computers and people, remained relatively fixed. The same is true of statistics compiled for the corporate network of Bellcore.

From such published data we suggest the hypothesis that the natural growth of data traffic in networks is exponential, and that the annual expansion is of the order of 100%. This reflects the intuitive understanding of communications managers that network traffic grows exponentially in a network with available capacity, and that the traffic growth only decreases when the network approaches saturation, discouraging users from further expansion of their usage.

The important conclusion here is that the natural growth is exponential, and that *at any given time we do not know the future uses that will continue the exponential pattern.*

For example in the last few years the popularity of the client/server architecture has fostered the growth of distributed computing, and the emergence of multimedia has begun to cause bandwidth demands to escalate. The sudden popularity of the World Wide Web and the Mosaic

hypertext user interface has begun a new underlying exponential thrust on Internet traffic, just as older interfaces seem to be reaching saturation. There is no reason to think that similar new traffic stimulants will not continue to appear in the future.

It is also important to say that this strong exponential growth is characteristic only of data traffic—not of traditional voice traffic, which grows only at about 6% per year. Thus the amount of Air Force traffic which is analog voice would not be expected to grow with time.

The value of a network increases nonlinearly with the number of users. This fact seems intuitively obvious when considered. Every user has his value increased by being able to access every other connected user. The more people you can communicate with, the more value you have. However, this proposition is often overlooked in the consideration of network security, where the idea is often to communicate with as few people as possible. These factors must be traded off, with explicit consideration being given to the value of increased communication connectivity.

2.4 Pricing of Commercial Communications Facilities

The price of commercial communications has decreased over time, both because of technology and because of emerging competition. Moreover, the price of capacity on a per-bit basis decreases with the bandwidth being purchased. Between these two factors, it is possible to effect an exponential expansion of communications capacity over time at a fixed price.

Historically in the United States the pricing of communications has decreased annually at approximately an 11% rate. On bandwidth, the general rule of thumb is that price per bit decreases as the square root of bandwidth. More accurately, the data suggests a guideline of price being proportional to the 0.6 power of the bandwidth.

Putting these two factors together, it is evident that twice as much capacity can be purchased for about 1.5 times the initial price, and that after about three years the price will have dropped about this amount. Thus without spending any more money, capacity can be doubled about every three years.

This exponential growth at fixed price is an attribute of commercial facilities that might not be shared by Air Force facilities, such as military satellite capacity. In the latter case it is not evident that either changing technology or emerging competition would serve to decrease prices so dramatically over time. This serves as another argument for the Air Force to ride the commercial market wherever possible in the purchase of long haul communications capacity.

2.5 Commercial Fiber Technology, 1994-2000

Commercial fiber technology provides an already powerful communications medium, and one that is growing in reach and bandwidth. Today's standard transmission systems send 2.5 Gbps per fiber on a single wavelength. Technology is rapidly enabling further exploitation of the intrinsic fiber bandwidth, which is perhaps 1000 times greater than that utilized currently. Wavelength multiplexing enables more channels to coexist with a fiber, while higher speed electronics enables higher speed transmission on each channel. By the year 2000, systems will probably use 10 Gbps as the standard speed per channel, with wavelength multiplexing of 4 or

8 channels on a fiber. Erbium-doped fiber amplifiers make upgrades of existing systems easy, since only distant transmitters and receivers have to be changed in order to increase capacities.

The bottom line of all this technological progress is that fiber is providing the cheapest medium for broadband communications, and that in the future it will become even more inexpensive on a per-bit basis. There is no question that fiber will be absolutely the least expensive way to get information from point A to point B. Therefore, when looking to future long haul communications alternatives for the USAF, commercial fiber must receive first consideration. Other media should be used only when fiber is infeasible.

2.6 Fiber Deployment, 1994-2000

The expansion of fiber deployment throughout the world is extremely rapid. In 1992 the rate of fiber expansion within CONUS was at the speed of sound (total fiber miles divided by seconds in the year)! However, the fiber backbones across the US are temporarily overbuilt, and attention has turned to other parts of the world, including in particular Asia.

In the near future it is likely that every city in the world with a population of 100,000 will have access to fiber. Furthermore, every land mass on which the Air Force might have to operate will be connected by fiber. This includes every island where there are enough people to have a conflict.

In spite of the overwhelming proliferation of fiber, it appears likely that there will be deep communications "holes" that still exist in the year 2000. These holes will be principally in central Africa and South America. Thus it cannot be said that the world will have been glassed over, and fiber access cannot be assured everywhere.

2.7 Fiber and the USAF

Because of the bandwidth and economy of fiber, it will be the medium of choice for digital communications. Furthermore, commercial computing technology will evolve based upon the availability of fiber communications. Thus the USAF should plan to use fiber long haul systems wherever possible. To do otherwise would cut the USAF off from the technology and economy of commercial information systems.

Even though fiber will have reached every population node on earth, it is not evident that the forward locations that the Air Force must reach will have fiber access. Any such location might possibly be on the order of 100 miles from a fiber route, but these miles might pose an insurmountable barrier. They might be spanned by a satellite link, but once satellite has to be used for any link, the satellites might as well do the whole job—the point of fiber becomes moot.

It is possible to bridge the geographic gap to a fiber terminus by setting up a portable microwave link. There are occasions where this might be a useful strategy during the middle phase of a deployment, after satellite has served for the initial deployment into a "desert" environment.

Unfortunately, even given physical access to a fiber terminus, there are other problems. One is that the fiber access will be controlled by local governments. It might be necessary to

reach a fiber terminus in an adjacent country. Furthermore, the fiber trunk might consist of a single line that could be disrupted or exploited by the enemy.

Thus, even though fiber must always be the medium of choice, it will not always be feasible to use. This presents a problem for the USAF, because their information usage and capabilities will naturally evolve with commercial technology based on expanding bandwidth. It will be necessary for the Air Force also to have information systems that use diminished bandwidth, and for preparation and training with these lesser capabilities.

2.8 Wireless

Some experts predict that as much as 80% of telecommunications in the year 2000 will use wireless access. Others, more pessimistic, say only about 30%. Either way, it will be a major factor in how telecommunications evolve. The following generations of wireless access are succeeding each other:

- Analog cellular (AMPS)
- Digital cellular
- PCS (personal, low power transceivers)

There will be a number of tiers of wireless access. Several global satellite-based telephone networks will be in existence in the year 2000. Iridium and Globalstar are contenders, but so are a number of other alliances. It seems certain that some of them will survive, while others will not. The major question is whether the satellite service remains an expensive prestige item for the special business traveler or becomes a low cost mainstay of developing countries.

For the Air Force the existence of these low earth orbit (LEO) telephone networks poses problems. How much capacity do they have? What priority might the USAF get in the event of a conflict? Could the enemy be stopped from using them? For now the answers seem to be: not enough, not much, and no.

The alliance between Bill Gates and Craig McCaw to build a broadband, LEO network with about 600 satellites is especially interesting. They would have to dramatically change the economics of satellite technology. If anyone can do it, they can, but it is a difficult bet. If it existed, it would be most useful for military purposes, offering worldwide, broadband data access.

In addition to the PCS networks, there is currently an effort to build a PACS (personal access communication) access infrastructure, with simplified, high speed links connecting phones in buildings and in pedestrian areas to a multitude of low power base stations. Such technology might be typically provisioned in developing nations as an alternative to wired infrastructures. By the year 2000 it might be widely deployed, so that wireless access would exist in all but the most inaccessible areas of the earth.

The technology of wireless has been largely point-to-point or many-to-one until now. It is likely that much technology will be developed for wireless networking in the near future. Wireless LANs will provide instant infrastructure for quick provisioning of networks. Such capabilities will be obviously important to military operations.

Another tier of wireless will be the unlicensed spectrum used now by cordless phones. This technology will be much more widely used in the future. It is also quite possible that more spectrum will be available for "cooperating transceivers," rather than being rigidly assigned through the Federal Communications Commission or its foreign equivalents.

Telephone and cable companies are also now rushing to use LMDS and MMDS technology for broadband broadcast. As of this moment, these are seen as quick start ways of getting into video access. The LMDS, at 28 Ghz, conflicts with the proposed satellite systems, such as Teledesic, but provides a great deal of potential bandwidth, whereas the MMDS at around 2 Ghz is easier technically and politically.

2.9 Network Security and Reliability

Unbreakable codes will exist as *standards* throughout the world (encryption techniques that are computationally infeasible to decrypt). This is a fact that NSA and other government organizations would like to prevent or ignore. However, the knowledge is already widespread, and the development of standards, cheap coders, and transparent interfaces will inevitably follow. This is both good and bad news. The good news is that security will become something that the Air Force shares with the commercial world. Its problems will no longer be unique. In fact, the Air Force might seriously consider at some time in the future the use of commercial security equipment and practices, instead of relying upon military techniques.

The bad news is that the enemy's codes will be indecipherable. (Although other methods of exploitation will always exist.)

The Air Force will find some of the new techniques of cryptography useful, including digital signatures, integrity checks, authentication procedures, digital time stamping, and electronic money. Undoubtedly, one of the areas receiving work will be the authentication technology, where biometric techniques will become more prevalent.

Network integrity and security should be a high priority concern of the Air Force. The Air Force should know how to keep its own networks up and how to take down the enemy's networks. Both will be very sophisticated technologies, and based upon past history the offense and defense will be well matched. It seems likely that will continue into the future, with hacker techniques keeping fully abreast of every firewall put in their face. Even national networks will be vulnerable to disruption by knowledgeable people.

2.10 Internet

The Internet will be of prime importance in the global information infrastructure. In fact, it will probably *be* the global infrastructure. At the current rate of growth the number of users equals the world's population in the year 2003. However, some saturation will occur, and a reasonable estimate might be several hundred million users in the year 2000. All the earth except central Africa and South America will be heavily connected. The network itself will be diverse and robust. It will be hard to take down, both technically and politically. In a conflict both sides would make use of the network.

The number of people on Internet will be multiplied by their use of software agents or knowbots. Agents will watch for events, will combine services, and will effect transactions. They will have sets of privileges that will constrain their abilities to act.

The Internet will become a major conduit for business, and will carry electronic money in various forms. Thanks to new protocols, and a continuous expansion of backbone bandwidth, the Internet will carry an appreciable amount of real time traffic, including speech, video, and multimedia. Video mail may become a popular feature of the network.

A huge amount of information will be readily accessible on Internet. Not only digital libraries and documents of all sorts, but even arrays of sensors everywhere on earth will be on-line. You will be able to access the temperatures in rooms throughout a distant city, or monitor the traffic flow. Through actuators you will be able to effect complicated commands at a distance. In the distributed network the speed of light will become an important limitation.

People everywhere on earth will be empowered to broadcast on the Internet. It is likely that no happening will go undocumented on the Internet. In a conflict amateur video cameras will go on-line to broadcast through the Internet. Common personal computers will be powerful video servers.

A structure for information will evolve that gives powerful search capabilities. Browsers will be simulated, 3-dimensional point-of-view like some computer games of today. There will be execution environments embedded in HTML documents, so that actions may occur at a distance. This is a very important development in network use, which will have a major effect within a ten year horizon.

Every soldier might well attempt to make personal use of the Internet in a conflict, provided he or she could get access. If wireless infrastructure becomes relatively ubiquitous, this might be possible. A tremendous amount of information will be there for the taking—for either side.

3.0 Personal Computing

Introduction and Definition

“Personal Computing” is a category encompassing the information devices likely to succeed today’s personal computers as commonplaces in personal and business lives. Thus this is not a category defined by technology per se; rather it reflects the convergence of technology trends from other categories, expressed in commercial products. This category is essentially one of packaging—what novel forms will various components and software be combined into?

In short, this category represents the trajectory of mass market information tools over the forecast period. Understanding this trajectory allows for exploitation of the economies of scale created by commercial deployment—including lowered cost, increased reliability and expanded understandings of systems performance. For example, if Pentagon planners a decade ago knew what Sun, Apple and SGI would be building today, how would that have affected planning? However, such exploitation is not without risk. Incorporation of complex commercially available systems may expose users to unknown risks of penetration and reliability failures.

Finally, it is vital to acknowledge the crucial role played by different age cohorts in defining this sector. The first PCs were invented by individuals who grew up in a mainframe time-sharing world, and thus it is no surprise that the first PCs (e.g., DOS) were near-perfect simulations of what a timesharing experience would be like at a local level. Now, a new generation raised on PCs is about to enter the workforce with ideas of its own. They are certain to redefine the PC sector as profoundly into something utterly new.

3.1 On the Nature of Consumer Electronic Diffusion

Success in this sector always looks like an S-curve—with a long anticipatory tail leading to a sudden inflection point. Figure 2 depicts exactly this pattern for the consumer devices that have insinuated themselves into our lives over the last few decades. This pattern offers important lessons for technology planners. First, do not overestimate the speed a new technology will go from initial introduction to consumer acceptance—the duration will more likely be decades than years. But once acceptance begins, expect it to occur vastly more rapidly than the apparent lack of earlier diffusion might imply. Fax machines took between 10 and 70 years to reach takeoff depending on how one measures the initial innovation, but once take-off was reached, fax machine penetration in US businesses went from less than 5 percent to over 70 percent in less than two years. It is a safe bet that the same pattern will be repeated by a host of new technologies in this sector over the next several decades.

3.2 Three Waves of Key Enabling Technologies: Microprocessor (1970 -95), Laser (1995- 2005)—and Sensors (2005-)

The primary factors shaping change in this category are social and cultural. Technology plays an essential, but secondary role, merely enabling the possibility of novel devices satisfying human needs and desires. That said, the technologies waiting in the wings will enable some impressive advances in the forecast period. The first decade of the personal computer was very much shaped by the availability of microprocessors cheap enough to afford one on our desks.




Figure 2 Technology Acceptance - S-Curves

The consequence was a 1980's decade preoccupied with *processing*. We processed everything we could get our hands on—words, images, figures, even video.

Now however, we are experiencing diminishing returns on the addition of processing power to processing-era tasks, and a second wave technology—the communications laser—has replaced the microprocessor as the key enabling technology shaping the personal computer industry. Just as the microprocessor slipping into our lives in the early 1980s hidden in PCs and other gizmos, lasers are slipping into our lives today hidden in everything from CD players to long-distance fiber-optic phone lines. If the 1980s represented a processing revolution, the next 10 years will be above all an *access* revolution in which the devices in this sector will be defined by what they connect users to. This access demand will in turn increase the demand for ever more powerful processors to support increasingly sophisticated access tasks.

Further out—beginning perhaps around 2005—a third wave is waiting in the wings. This third wave will be defined by the advent of inexpensive high-performance sensor technologies. This will trigger a decade preoccupied with connecting our personal computing devices to the real world—giving them the ability to sense objects and events around them. MEMS (micro-electro-mechanical system, chips that combine mechanical functions with electronic functions) technology is an early harbinger of this third wave, but it is by no means the only technology enabling the sensor revolution.

Note that all three technologies appear in all three stages; but at each stage, one technology takes the lead in setting the applications stage. The effect among advances on all three fronts is synergistic. In the 1980s, advances in processing made early use of lasers possible. Now, advances in laser-enabled bandwidth are creating new demand for ever more powerful lasers to, for example, process video images sent down phone lines. And the advent of cheap

sensors will have similar impacts on creating demand for processing and communications between sensor and the ultimate user.

3.3 The Fragmenting PC Industry

The first casualty of the shift from microprocessor to laser is the 1980s-vintage stand-alone PC, now quickly headed for the technological scrap heap because its processing-driven nature leaves it unable to cope with an access-driven world. A "personal computer" is exactly what its name implies, a device designed to stand alone, connected to little else than an AC outlet and, occasionally, a phone line. It is above all a processing engine, into which most information goes in by keyboard, gets cleaned up on screen, and eventually emerges as a paper document from a printer. One can hook a 1980's vintage PC to access-era conduits, like local area networks or remote information sources, but it will quickly begin to gasp and stall under the unfamiliar load.


In fact, what once was the PC market is splitting into three new categories: A new class of device on our desks, novel entertainment systems in our homes, and a third class of ultra portable gizmos that we will carry in our cars, briefcases and even on our persons. A few of these devices will outwardly resemble PCs, but in fact they will be as different from the devices of the 1980s, as today's cars are compared to the original horseless carriages. And for the next decade at least, all will share a common denominator—they will all be access tools, defined by what they connect us to.

3.3.1 Workstations Will Replace PCs on Our Desks

The diffusion of cheap laser technology has already triggered a shift from standalone PC to highly connected workstation. Our new desktop devices superficially resemble PCs, but their functions are access-driven from the start, defined above all by what they connect users to, not by what they process. These new tools will mature into "workstations" in the truest sense of the word, personal windows onto an ever growing ecology of information and computing devices. Some information will still be entered manually, but the bulk will arrive via LAN or external communications channels. Keyboards and mice will be tools for manipulation rather than entry. On the output side, less information will emerge from the printer slot and more will be dispatched in electronic form to other remote workstations and electronic display devices.

The world that workstations open onto is growing explosively, thanks to laser-enabled bandwidth and storage economies. In turn, these new fiber digital highways are encouraging the diffusion of secondary transport modes such as cable and wireless conduits, the surface roads and alleyways of the access decade. The steady stream of surprises swirling off of the Internet is but the harbinger of bigger surprises to come for workstations as information windows. E-mail and Web-surfing are hot for the moment, but weirder exotica like multiple user dimensions (MUDs) and M-bone multicasting are waiting in the wings to further enrich the world our devices access.

The rapid diffusion of the Internet is a key enabling factor in this sector. (see Figure 3—Internet Growth.) The Internet has a double effect on diffusion for workstation technology. First, it is an excellent conduit for moving information between workstations, but second and more importantly, it is a technology very much accessible to individual entrepreneurs. It is this latter population that will deliver a steady stream of interesting and destabilizing surprises.



Source Network Wizards

Figure 3 Internet Growth

The explosion of Internet use will also likely yield new kinds of access devices. For example, well before 2000, we will see inexpensive, easy to use terminal devices that rely on advanced Internet technologies (e.g., Sun's Hot Java) to deliver powerful experiences without the need of massive amounts of local processing or storage. If manufactured in volume, these devices could be delivered for costs well below that of a French Minitel terminal.

3.3.2 "Crays on a Tray" In Our Homes

Laser-enabled bandwidth in our homes will also eventually lead to a new class of entertainment devices in our homes. These will grow out of the TV/VCR/ video game player complex that exists today. In effect, we will have a "Cray on a tray" in a box sitting next to our TV monitors, delivering a host of new entertainment services. The most prosaic applications will be movies on demand and simple home shopping. The big winners and surprises will lie elsewhere, in new kinds of group activities evolving out of activities like MUDs on the Internet today.

However, there are two important caveats affecting this sector. First, it will be the slowest of the three sectors to mature, for the simple reason that the deployment of the requisite networks (ideally, a two-way, fully-switched megabit conduit) is proceeding slowly, even in the United States: less than 10 percent of US households by the year 2000. Look for this sector to start coming into its own in the years following 2000, and diffusion of high-performance networked entertainment systems is likely to reach take-off around 2005.

Secondly, this entertainment device will not merge with desktop PCs in the home. There are two screens in our houses: one that we sit 8 feet away from, and watch while leaning backwards on a couch, and a second that we sit three feet away from, and lean forward while we use it. Today, we call the former a TV, and the latter a PC. While the details of these devices will change profoundly, we will still have 8-foot and 3-foot screens in our homes decades from now. And, we will probably have a third as well—screens on the successors to today's phones.

3.3.3 Information Appliances Will Become Fixtures in Our Lives

The third sector emerging from the dissolving PC market will be the far stranger world of "information appliances"—inexpensive, radically accessible, high-performance information tools utterly unlike the PCs we use today. The difference between PCs and information appliances lies first of all in their specialization of use. Today's personal computers are general purpose tools designed to run a wide range of applications, artifacts of an age when microprocessors were so expensive we could afford to have *only* one on our desks. As David Liddle of Interval Research notes, PCs are the equivalent of Swiss Army knives for knowledge workers, general purpose tools of last resort rather than special purpose tools of first choice. Who would actually use the saw on their knife if they had a real saw within arm's reach?

Information appliances will combine the information richness of personal computing with the low-cost and hardware elegance of consumer electronics. The first information appliances have already begun to appear: Apple's Newton, Sony's MagicLink, and the Sharp Zaurus are but the first tottering steps of these infant sectors. Ten years from now pen-based notetakers and email/fax communicators will hardly draw the attention of passers-by.

But this description belies the sheer chaos and confusion that will accompany the menagerie of information appliances lurking on the horizon. These new tools will come in every conceivable shape and size. Manufacturers will entice us to pack them in our briefcases, stuff them in our pockets, and toss them in our cars. We will even wear some as integral parts of our clothing, a new kind of "information exoskeleton." Many will be silly, faddish gizmos, while a minority will quietly become essential personal tools. But all will share two features in common—all will be intensely personal, and none will remotely resemble a PC.

The diffusion of these devices will be slower than it at first seems. There is no shortage of candidate devices today, but the vast majority are mildly useless curiosities. However within five years, the first useful versions will begin appearing, and a decade from now, we will see some very impressive pocket-sized tools. By then, the diffusion of these gizmos into the work place will create far more headaches for information managers than the PC did a decade ago. For example, unlike PCs a decade ago these gizmos will include built-in communications capabilities spanning everything from simple wire-based telephony to infrared, packet radio, and cellular voice and data. MIS managers despaired of early PCs because they *couldn't* communicate with their mainframe systems. Information appliances will present them with the opposite problem. Conditioned by experiences with desktop workstations and aware of the communications potential of their new portable tools, knowledge workers will be exasperated when companies can't make their information and communications systems interface with these appliances. Just as PC zealots once took matters into their own hands, the appliance owners of the 1990s will quickly learn to communicate on an ad hoc basis with everything from their

workstation to the fax machine and digital copier. An absence of universal operating and communications standards will only add to the confusion.

3.4 A Functional Categorization of Future PCs—by Size

Perhaps the single constant amidst all the turbulence in this category is human physiology—how much can one carry, what are optimum viewing distances, etc. The consequence is that size is a powerful tool for sorting what is certain to be a steadily expanding ecology of gizmos of all shapes and sizes. Here is one sorting:

3.4.1 Less than 10cm (Ounce-Scale Weight) = Wearable Computers

This is an all but unrealized category today, inhabited by a few pagers and some electronic jewelry novelties, and of course, digital watches. These are devices small enough to fit in a pocket or be pinned to clothing. Too small for anything but a minimalist screen, but acceptable for sound i/o. Look for short-hop wireless gizmos in this category that communicate with larger personal devices. For example, pocket dictating machine that uploads speech via wireless or infra-red to PDA or larger machine. This category is unlikely to begin really taking off for another decade, until then remaining a zone of mostly novelties with a few practical exceptions.

3.4.2 *10cm - .25m (up to 1.5 pounds) = PDAs

These devices are as large as book-sized, and book weight, though a video cartridge represents the center of the size/weight form factor of the this zone. This category is already populated with first generation devices (e.g., Newton, MagicLink, Sharp Zaurus, HP 100), and for better or worse, has been dubbed by the industry as PDAs—“Personal Digital Assistants.” Second wave PDAs will arrive in the near term (next two years), and the category will reinvent itself every 2-3 years for the next decade.

3.4.3 *.25m - .5m (up to 5 pounds) = Laptop PCs

This is a well populated category today, moreover the form factor (keyboard, screen, supplemental pointing and input devices) is quite stable, and likely to remain so for the forecast period. However, the components, functions and operating systems will evolve radically over the forecast period. Within 10 years, “portable” will mean under two pounds, and with a form factor that one can carry the device as well as an armful of other items as well.

3.4.4 *.5m - 1m (over 5 pounds) = Desktop PCs

Like laptops, this is also a stable category: screen and keyboard, plus assorted i/o peripherals. And, like laptops, it is a category that will see dramatic evolution, though very much along a continuum. We have already seen this category absorb a profound shift in function from standalone processing engine to highly connected access tool. The access function will now dominate through the balance of the forecast period, but processing will also play a crucial role. Look for major operating system rollover every 8-10 years.

3.4.5 * Over 1m (over 5 pounds) = Entertainment Devices

This is the zone of set-top boxes, TVs and related gizmos. This will be important in terms of exploiting economies of scale. Air Force pilots are already doing situational awareness

training on desktop PC systems employing software that is only modestly modified from commercial versions. This zone could provide far richer training opportunities.

3.5 Factors Retarding/Accelerating the Evolution of PC Sector

A variety of factors are at work that are likely to accelerate or retard advances in this sector. They are summarized below:

3.5.1 Batteries

Best estimates are for no more than 20% improvement in battery performance over next decade. Improvements must come in combination of battery performance and low-power processing. Xerox Parc researcher Mark Weiser has an interesting view of this, thinking in terms of mips/joule (given a D cell = 2^{16} joules): Pentium: 100mips/25 watts=4 mips/joule, 1/2 speed 3.3 volt Pentium: 50 mips/ 5 watts = 10 mips/joule, Motorola 68010: 1 mip/0.5 watts = 1/2 mips/joule. This is an area worth a close look: an unexpected advance in battery performance could have a dramatic impact on the information appliance sector in particular. The consequence is more than mere device life: more power translates into more processing cycles and thus more powerful functionality. By 2005, the average consumer could carry the processing equivalent of a Cray in their pocket.

3.5.2 Basic Research

Product innovation in the PC sector is highly dependent upon a steady stream of enabling technologies. Until now, many of these technologies have emerged from defense-related R&D—for example, CCDs currently used in consumer video cameras. As defense-related research wanes, an increasing portion of this work is being undertaken by private corporations (e.g., Intel and the development of the Pentium and P6 chips). But it is an open question whether this private, highly directed research can effectively substitute for more broad-based government-sponsored research that might lead to far greater innovations and surprises. If defense-related research—indeed, government-sponsored research in general—declines, which stakeholders will take over the basic research “pump-priming” so essential for innovation in this arena. Will the research enabling the devices of 2010 be done at all in 1997?

3.5.3 Competitive Quest for Novelty

The pressure on mass-market manufacturers to come up with new products is enormous. The consequence will be a continuous effort to take some newly available piece of technology and use it as the cornerstone for a new product or product category. In addition, we are likely to see the creation of entirely new industries on a decade-scale basis. Just as the personal computer industry appeared out of nowhere in 1980, look for new industries to emerge around sensor-centric systems and autonomous systems in the decades ahead.

3.6 Visions/Discontinuities and Implications

This is an area that is likely to create discontinuities even as it realizes earlier visions. For example, it is in this sector that Vannevar Bush's 50-year old vision of the Memex is likely to take the form of a real product. This section offers a speculative exploration of possibilities.

3.6.1 Information Exoskeleta

We end up wearing a small ecology of information devices on or near our persons—gizmos in our pockets, on our wrists, in our briefcases. These devices increasingly intercommunicate via a personal network and serve multiple access and processing functions. Business users in particular are heavy users. This is the civilian equivalent of the Army's wearable infantry information systems now under development.

3.6.2 Non-Interactive (Autonomous) Computing

The current information systems paradigm makes human/computer interaction the main event, and this category's title assumes that it will remain the main event. Perhaps more machines will operate without human interaction on the behalf of humans, than will interact with humans directly. Will our homes be run by a small community of devices constantly in touch with each other, but which never directly interact with the humans they serve?

3.6.3 Consumer Infobots

Will we see cheap, small infobots become a reality in the forecast period? These infobots could become alter-egos on networks, hunting down information, conducting transactions, and monitoring critical pieces of data. (See Intelligent Software Agents, Chapter 5.)

3.6.4 MUDs

MUDs—Multiple User Dimensions—are teetering on the verge of becoming a major consumer event in the next two to three years. Email will seem quaintly un-hip. Within the decade, look for consumer systems optimized for MUD-like interactions. MUDs are covered in greater detail in Groupware, Chapter 14.

3.6.5 Civilian War-Fighting

What seems exotic to the Pentagon today will likely become a consumer hobbyist attraction within a decade or two. Transform today's combat virtual reality (VR) networks into the commercial sector, and you have the biggest theme park in cyberspace, with 15 year-olds flying fighter simulators in battles against others of their age cohort. We note that the developers of a leading fighter simulation in 1995 have purposely fuzzed and degraded performance on the next version to be released, for fear of giving Third-World aviators a powerful and inexpensive simulation that could put US pilots at a disadvantage.

3.6.6 The End of WIMP Interfaces

WIMP (Windows, Icons, Mouse, Pull-down menus) interfaces have become the norm in the PC sector over the last decade. We are likely to see their survival for the next decade, particularly given the rapid acceptance of Windows 95. However, it is critical to keep in mind that this interface metaphor is already over 20 years old, and is merely one of many possible ways of interacting with computers. WIMP interfaces won't disappear, but they are likely to take a back seat to newer metaphors over the next three decades.

4.0 Human-Computer Interaction

Introduction and Definition

As soon as the users began to be other than the designers of computing machinery, the so-called human-computer interface became an object of design itself. That boundary has been loosely referred to as an "interface", much as any other computer peripheral. But the term "interface" is losing its hold, mainly because it suggests a sterile, inflexible specification that doesn't match the steady introduction of more abstract and natural means of communications we will call modalities. The more descriptive and anticipatory term is "interaction".

The next twenty years will see two important ways computers will come to serve us. The first and undoubtedly most compelling vision is to see their migration from a hands-intensive tool toward a delegatable assistant. This concept is not an insensitive attempt to gloss over difficult concepts with anthropomorphism but to point to an inevitable direction whose ultimate measure will never quite be realized. Simply put, the evolution of human-computer interaction should be thought of as a direction and not an end. Straining to grant the human user an interaction for which training is not a prerequisite, the ever-affordable surplus of computing cycles will be more and more dedicated to making the computer not just user-friendly but user-like.

Thus, human-computer interaction (HCI) encompasses a very wide range of functionality and work. In general, it is all the functions, mechanisms, and conventions that provide the means for users to interact with computers. Some interface specifications or definitional languages may exist until natural language dominates. In a research context, it is all the enablements whereby humans and computers interact and is NOT an interface in the sense of the specification of a standard.

4.1 Principal Motivations

As with most machines man has built, the computer requires its user to learn what functionality the machine has as well as how to interact with it. Though the potential for vastly easing these burdens has been there in the case of the computer, it has given little quarter to its operator....so far. But, as stated above, the computer has the potential to offer its user a more abstract way to convey orders and the next decade or two will see substantial changes and improvements in this regard. Table I lists some of the reasons why this trend will occur.

The motivations listed in Table 1 can be aggregated into just two: the computer as an increasingly delegatable assistant to an information-centered user and the computer as an augmentation to the physical, not just intellectual human.

The more straightforward use of computers today is as a framework within which application programs are run. The processors, the connectivity, the operating systems are becoming invisible...subordinated to the applications programs at hand. Within a decade those too will begin to sink below the conscious surface, yielding to task-oriented computing at the next level of work abstraction. To make this possible, the way people interact with computers will take on some of the attributes used when they interact with other people. Within ten years, very acceptable speech recognition and understanding software will permit a wide range of computer tasking now confined to the keyboard and mouse. This will be important wherever keyboards are

Table I Motivations for More Natural Human-Computer Interaction

- The present wide range of inconsistent interaction environments requiring retraining as a user moves between systems
- The need for easier, more intuitive interactions to achieve objectives making human language the dominant, but not only, interaction language
- Tasks of greater complexity can be accomplished using higher-level abstractions with attributes such as completeness, veracity, timeliness....
- The notion of delegation - the computer or computer-based machine accepts increasingly abstract instructions and tasks
- The computer as a sensory and action mediator, matching human dexterity and senses to situations in which humans cannot function or easily adapt.

impossible to accommodate. This change does not mean over that time frame the computer will necessarily display very much reasoning power, but it should be capable of both understanding and synthesizing natural language.

An important conceptual aspect of HCI is just what the human is interacting with. Today's platforms suggest that the user interact with a local client or host; whatever is personal to the user is located within easy, direct reach. That notion may broaden considerably toward a personal space that is quite virtual.... distributed in ways the user may be unaware of. The virtues of such a convention may be that the user's location may one day be totally irrelevant. Present systems hint of this now but all access is cross-net to the user's personal physical machine.

As mentioned, the second compelling use of computers will be to augment the physical rather than just the intellectual human. While the range of such use may not be as broad as more mind-centered work, there are many situations where our dimensionality, our sensory package, our reaction times, our motor sensitivities, or simply our remoteness are not adequate to the task. Through relatively new concepts such as telepresence, computers will effectively transport us into other dimensions and places to let us view or interact with what we find. (See Personal Computing, Chapter 3.)

4.2 Important Counter-Pressures

Undoubtedly the most important resistance to a common and consistent interaction environment is the progress of technology. That movement is apparently in the direction of more natural, more abstract, more consistent levels of interaction but its form is of a few companies vying for the user's devotion to their operating environment. One important aspect of that competition is that the progress in functionality within each camp is both small and cloistered enough to keep users captive, all the while forcing a lifelong user adaptation. Today's user is a paragon of tolerance to these parallel and lucrative strategies of technology advancement and payment for upgrades. The concomitant advancement of the technology and this tolerance for small, incremental change are impediments to a global and consistent standard.

Another obstacle against an HCI standard is the continuing change in the nature of computers. In this case we refer mainly to their shrinking size and, to a lesser degree, their ubiquity. As computers get smaller, the keyboard is no longer an option and, if one is to retain some fluency in the interaction, speech understanding becomes important. Total ubiquity of computers is the preferred future for some, with the interface becoming either unobtrusive or a constant set of simple functions, the first implying a more natural interaction. But ubiquity may also arrive in the form of the long-discussed information utility with its countless uses. Managing the access and those uses would benefit from an intuitive and natural HCI. Of course, ubiquity also means that many users will only see the consequences of computers with no direct accessibility to them at all.

Environmental factors can also compromise or even invalidate a common and consistent interaction environment. In the wireless world, high noise or multipath will force more narrow bandwidths that may cause abbreviated frameworks for interaction, particularly if the HCI is centered in some remote server.

4.3 Models and Modalities

In this HCI discussion, "model" refers to the range of computer types and, to some extent, how they are used. "Modality" refers to the channel or medium that is used, such as pointing, speech, or keyboard. Table II shows a taxonomy of different models for HCI.

Table II Important Types of Human-Computer Interaction

- Desktop Model - Portability is not primary. Wired media is expected. Large displays, rich input output modalities.
- Mobility Model - Portability and power saving are paramount. Radio or wire. Much more constrained input output modalities.
- Distributed Model - May have some of each of the above. User sees constant environment regardless of location.
- Virtual Reality Model - Immersive interaction within a totally simulated environment.
- Augmented (Enhanced) Reality Model - Simulation overlay to real world environment. Generally a real-time process.
- Telepresence Model - Computer-mediated interaction with a real environment. Generally a real-time process.

Apart from the *models* in Table II is a collection of underlying *methods* used as a basis for building the interaction environment. Examples of these are: situation theory, tasking models, agents, learning-metaphors, requirements/tasks/usability, tool-based, common ground, and black-boards.

The spectrum of interaction modalities is defined by our communications-oriented senses: displays, speech, handwriting and gesturing, text, haptic feedback, and their integration. In most tasks we undertake, we benefit from the use of multiple modalities. Computers to date have been dominated by the keyboard, display, and mouse (KDM) convention. But the recognition accuracy of speech and handwriting are opening the door to other modalities...albeit slowly. Multiple-agent systems are now being demonstrated that permit very natural integration of the different modalities common to HCI.

4.4 Realizations Within 10 Years

In the context of the models mentioned above, here are some predictions about HCI over the next decade:

- The Desktop Model - As the best example of how the need for standardization defeats rational design, the omnipresent QWERTY keyboard seems eternal. Because of the need for authoring text, the keyboard will remain. But because speech is the most convenient method for addressing or seeking objects not displayed, it will see increased use and in circumscribed environments will yield almost 100 percent accuracy. Modality integration will be greatest under this model and expect a few interaction environments to illustrate modality independence. The infrastructure will be present to reach 100 Mbps peak data rates at the desktop terminal for about 25 percent of the market.
- The Mobility Model - Computing power in laptops and PDAs will be roughly equivalent to desktops but the input/output methods will lag. Speech and handwriting recognition will improve and some will be integrated together. The mobility model will be dominated by lap-tops of ever increasing power. Some smaller ultra-portable wireless systems will use speech in addition to a stylus and buttons.
- The Virtual Reality Model - A broad assortment of virtual reality systems and their components will become available through the entertainment industry. The present major impediment to VR, the combination of display resolution and rendering speed, will improve greatly, offering much more realistic, totally synthesized environments. But matching the eye's capability using head-mounted displays is a very daunting task. Extremely high resolution is needed and the costs to achieve even a million-pixel fields is now prohibitive. Today's displays also have a fixed focus for all objects, leaving depth perception and occlusion unreconciled. This and several other factors may lead to VR sickness after long use for some users.
- The Augmented Reality Model - Augmented reality systems require more modest computing power but a great deal more registration accuracy, probably at least an order of magnitude. This need is in the head tracking system, with the amount depending upon the size and distance of the object being interacted with. This is essentially the spatial synchronization problem between the real and synthesized worlds. The applications such as training, maintenance, and medical care are important enough, however, that the AF should help leverage the development of

this technology. Typical will be the projection onto real objects (or on special eyewear) of graphical images that reflect some training aids to the user. Some systems using this will be ready for AF employment within 10 years but it is not clear whether the non-DoD market will be big enough to do this without military investment.

- Telepresence Model - Systems will emerge for medicine, hazardous waste handling, and other tasks with accessibility problems. Telepresence surgical systems will be in clinical trials or beyond. Telepresence is probably the most relevant information technology for RPV systems, including the possible extraction of the fighter pilot from the cockpit. Scale transformations are also a use of telepresence. Satellite repair or space station construction are examples as are microsurgery and microsystem assembly. Telepresence systems will permit distance-spanning meetings with far greater realism than the "head in the box" video or computer conferencing of today. Again, AF R&D investment will likely be needed to speed telepresence technology along.

Perhaps the most intriguing question for a 10-year projection is just how content users will be to keep using the present dominant keyboard-display-mouse (KDM) configuration for HCI. It is clear it will not vanish and may have considerable longevity. The information technology world has other good examples of inefficient or awkward conventions that are maintained simply for the sake of standards and often because they are implemented using products that have very low marginal cost. KDM will likely persist.

KDM dominance will just slow but not prevent the inevitable movement toward interactions that are more natural for humans. It may mean that new HCI functionality may have to emerge from niche markets. For example, speech recognition, arguably necessary for input in powerful but small computers, has found its first utility in the labor cost saving replacement of operators. Continuous, speaker-independent, large vocabulary, real-time, domain- or context-constrained speech recognition is beginning to find utility. Speaker independent accuracies have gone from about 21 percent word error rate in 1987 to 3.6 percent in 1991 for a 1000 word Navy battle management vocabulary. More recently, error rates on Wall Street Journal text with much greater complexity were about 7 percent. Couple such recognition accuracy with the added utility of circumscribed, context-aware, natural language understanding and very useful applications arise. Thus, the increasingly capable ability to understand speech will find more and more applications in HCI over the decade.

Another evidence of progress in HCI is the integration of the important interaction modalities into more natural and intuitive combinations. Speech understanding along with the pen-based inputs of text, symbols, and gestures are being integrated through the use of agents that not only enable better interaction but are adaptable to both task and user preferences. An example is the Open Agent Architecture, a schematic of which is shown in Figure 4.

Our access to broad information stores has been much simplified by the search and relational addressing conventions of, say, the WWW. But the fusion of data taken from disparate data bases ultimately needs knowledge mediation and this decade will see research demonstrations of that integration. (See Information Access Technology, Chapter 6.)

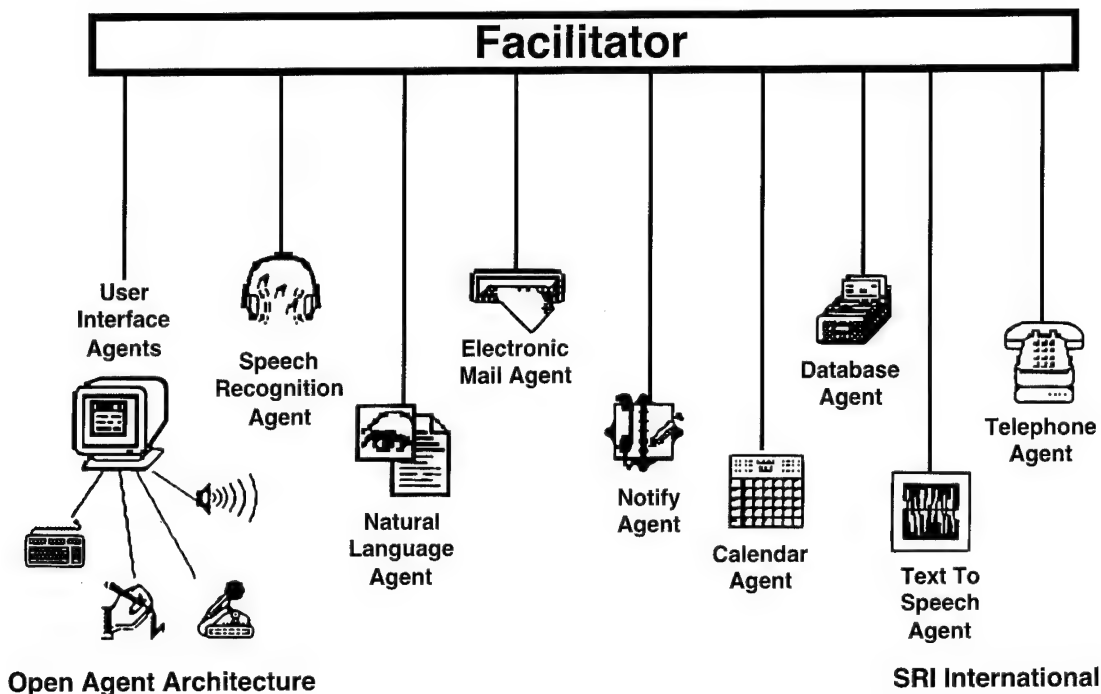


Figure 4 SRI's Open Agent Architecture

4.5 Realizations Within 20 Years

Again repeating the model metaphor from above, here are some longer term projections:

- The Desktop Model - There will be a very wide variety of task-oriented versions of HCI. Interactions between user and machine will be in natural-language, semantically flexible, and redundant. Displays and scanners will operate consistent with the limits of human vision. From a "home" machine full networking will enable direct bilaterally safe access to other machines, most with a common "look and feel."
- The Mobility Model - This will be embodied in three versions: ultra-portable or wearable, palm-tops of pocket size, and notebooks. The first is for speech input/output only or, in some cases speech/pen-based input. Some network connections and high resolution, eye-worn displays that can also double as projecting devices. The second is now called PDAs. They will become communications-rich terminals with, in some instances, enough local storage to make the limiting propagation and noise environment adequate for communication via differentials; i.e., only information different from that locally stored. Notebooks will have computing power equivalent to most desktops with perhaps less storage and I/O capabilities. (See Personal Computing, Chapter 3)

- **Virtual and Augmented Reality Models** - Improvements in these areas will be incremental over the second decade but that will ensure, at least technically, the availability of high resolution synthetic environments from virtually any location. Collaborative virtual and augmented reality environments will also be distributed and available. Adequate resolution displays and tracking system accuracy should be available.
- **Telepresence Models** - This type of computing will grant increasing accessibility for humans to otherwise inaccessible space. In medicine computers will mediate delay-bound remote surgery, microsurgery, all types of minimally invasive surgery and will do so giving the surgeon consistent, natural, sensory-rich accessibility. Manipulation of micron-sized objects with synthetic haptic feedback will enable the assemblage of nanosystems. Imaging and other human-compatible sensors will be compressed over communication links to permit remote experiences of high realism. Teleoperation will also benefit from computer-mediation that can, for example, guarantee the avoidance of undesirable actions.

Other supporting technologies will also be present:

- Totally integrated modalities will offer the flexible means to express a given semantic notion.
- Intelligent agents will learn users' intentions and either advise or react accordingly.
- User-crafted rules, conventions, appearance will yield a "behavior" for a computer keyed to its present user. That tailoring might then appear at any access point.

Displays will be both high resolution and very large. Mobile systems will have projectable displays.

4.6 Untethered Realizations

- **Human interaction** will be defined only by the task and user preferences - Until now the user has been more adaptation-prone than the machine. That will reverse itself. All tools should be selected or constructed by their user(s) according to the task. The computer as a tool is so malleable its utility seems boundless.
- **The transparent computer** - The computer is the first machine that can take tasking and delegation in abstract, human-compatible terms. How it accomplishes the task can become totally transparent (not necessarily irrelevant) to the one who is using it. One may wish to give it form, but the input and output devices may also be hidden and ubiquitous. If the interaction involves physical input/output objects, then the user will obviously face associated location constraints. (See Artificial Intelligence, Chapter 7.)

5.0 Intelligent Software Agents

Introduction and Definition

The notion of software agents has been with us for a long time. Certainly since the formative years of artificial intelligence, designers have flirted with the notion computers might one day have a degree of apparent independence that would warrant their being called an agent of the user. But imbuing a machine with the attributes of an agent has been difficult to achieve. Not until processor power began to exceed appreciably that needed by the application at hand, have resources been available locally to give the user this type of support. Now that has changed. With evermore affordable processing power available at the user's first point of interaction to this new distributed computing world, there is the luxury of giving the machine attributes that, to the user, seem autonomous and intelligent. Whether they are either, of course, is in the eye of the beholder. Semantically, an agent is simply one authorized to work on behalf of or as a representative of another. In computer science research, one expects an "intelligent" software agent to use reasoning and persistence in performing its assigned task. Other, more human-like attributes such as trust are more controversial.

The lack of a precise definition of a software agent, unfortunately, gives the developer and marketer wide latitude in just how much or how little functionality is present. At their best, software agents are capable of representing a user or owner in the accomplishment of specified tasks without his or her having to prescribe or even be aware of how it is to be done. That degree of detachment has some influence on whether the user considers the agent intelligent or not. The more computer-aware the user, the more reasoning power and autonomy the agent must have to be termed intelligent.

In the development of agents there may be a propensity to ascribe human attributes to a program as its functionality increases. Because computers are the first machines that can take on tasking through abstract, human-compatible language, there is some reason for this tendency. But care must be given in not misleading users as to the reasoning power and adaptability software programs like agents actually have. It is doubtful, however, enough precision will emerge in describing such capabilities, so imputing functionality not actually present will be an ongoing problem. Because of this difficulty, the term agent has mixed acceptance by many experts. That reticence will not likely prevent it from being more popularized.

5.1 Principal Motivations

Why are software agents important in the evolution of computing? Listed in Table III are some of the reasons why agents will be of increasing importance in dealing with a complex world of distributed information and resources the normal user will find confusing and threatening.

Table III Motivations for the Use of Software Agents

- The quantity of information will be too vast and its quality too uneven for most humans to suffer through.
- The locations of desired information are too broad and nonintuitive.
- An increased "flatness" (lack of hierarchy) in the world of sources and sinks for information makes dealing with it less understandable. This gets emphasized in future peer-to-peer systems.
- Data and database heterogeneity demand a variety of translators.
- Stored media will have increasing dimensionality in various formats.
- The notion of delegation, broadly defined, will become more available.
- A broadly accepted commercial infrastructure for on-line information will provide a more consistent interface.
- The need to hide complexity.
- The need for new programming models for a distributed computing environment.
- To need for improved human-computer interaction.

5.2 Important Counter-Pressures

Even though the above motivations seem to make agents inevitable, there are near term obstacles. Some are due to the difficulty of the technology itself while some are due to the difficulty in vendors or users being able or wanting to organize themselves.

Table IV Obstacles in the Development and Use of Software Agents

- A perception that agents pose a risk as they masquerade an owner's malintent.
- The lack of a formal framework for a trusted agent. Will it perform as specified?
- "Training" agents is difficult. Can they be trained to do only prescribed tasks in a remote setting the user may not know or understand?
- The lack of a common language - Too many conventions, not enough translators (scripting and interpreter languages are the most favored at present).
- The absence of commonly accepted models of reasoning or negotiating.

5.3 Functionalities, Types, and Models

Agents are clearly enjoying wide discussion in the *user community*. Given that and the above background, it is perhaps best to begin a discussion of agents using some functional examples. If the popularity of the concept continues, there will be countless versions of them, but here only four will be considered:

- Advisory agents - Those agents able to monitor a situation and give feedback with or without recommendations. Generally, application-specific. Monitoring, at some level of sophistication or abstraction, should be an attribute of all agents.

- Personal assistants - Most likely to appear as adjuncts to human-computer interaction (HCI). Will offer assistance in specifiable tasks.
- Traveling (Internet) agents - Roving, mission-specific, with broad awareness and interface potential. [Other than some disk access, most all web retrievals now run entirely on the user machine. Where information gathering processes run is ultimately a matter of money and risk.]
- Multiple Collaborating Agents - Multiple agents with some common goals, of varying sophistication; may be physically or logically separated.

Another, and perhaps more general way to describe intelligent agents is by their degree of intelligence. Discussing the amount of intelligence in anything carries both difficulty and controversy. If agents are an inevitable direction in software development, we must have a way to discuss them. Here are three types of intelligent agents:

- Directed-Action Agents - Has fixed goals but can react to the data and the environment it encounters as long as they have been explicitly anticipated. Little reasoning ability except for self recovery. Might be termed a “do-this” agent.
- Reasoned-Action Agents - Has fixed goals and is able to monitor other objects (data and processes). It can reason about what it encounters and take alternative action. Being able to reason implies that it has both a knowledge base (data and rules) and some processes to use. Might be termed an “achieve-this” agent.
- Learned-Action Agents - Has the above capabilities plus it can accept more general goals and is capable of altering or adding to them under guidelines. Has broad awareness of its environment, the data it encounters, and, importantly, itself. Might be termed an “accept-this” agent.

Whether or not software agents will catch on in the *development community* is in great measure determined by the underlying models on which they are constructed. Attributing intelligence to agents means they will likely be built on the foundations of AI. Agent models are often based on concepts such as perception/action, belief/desire/intent, expert systems, game theoretic, and others. AI programs are widely embedded in existing software and AI will continue to provide the foundations for new capabilities. The technology of distributed computing will also be a necessary component. Both are needed over the long term. (See Artificial Intelligence, Chapter 7.)

Lastly, agents will become so much a part of the HCI that it will be difficult to separate the two fields. Agents will almost certainly play a role in the next step in HCI as applications programs give way to a more integrated, task-oriented type of computer-based work. Another obvious need for agents, however, will be in powerful but ultra-portable computers (e.g., PDAs) where the more traditional input/output modalities such as keyboards are not available.

5.4 Realizations in Software Agents Within 10 Years

Over the next decade all of the above types of agents will be introduced and in common use at some level of sophistication. The rapid growth of the Internet may draw the most attention and thus the *traveling agent* may get broadly developed first. Agents that do appreciable

remote execution will be the most difficult to “host” or “serve”. The term itself suggests a warning to some but the acceptance of someone else’s code in your machine is already widely practiced under not-so-pointed terminology. WWW home pages, for example, are executable codes that the user invites in now. This is clearly an issue of security and therefore the initiative towards encryption and other solutions now underway on the Internet will continue with emphasis if business is to thrive there.

User or system assessment and characterization by a local *advisory agent* in a carefully circumscribed domain will be available. A good example of an advisory agent is its use in an instructional setting. Here an agent will start with a basic set of instructional goals, add an ability to monitor where the student is in relation to those goals, and then set the course of instruction issuing feedback, including recommendations, along the way. A natural evolution in the sophistication of situation assessment and how it is presented will make the agent appear more intelligent and interactive.

The *personal assistant agent* will hit the marketplace in 1996 with Apple’s new operating system. These agents will assist the user in specifiable tasks such as electronic mail, calendars, conferencing set-up, object search, and so on. These agents, unlike the simple macro-recording agents of HP’s New Wave, are reputed to have some ability to reason about what they are “seeing” and to act according to some prescribed user guidelines.

An important question concerning multiple *collaborating agents* is the level of complexity of a given agent. It is doubtful that small elemental agents can aggregate into something powerful or “intelligent.” So, the question is what should the atomic level of an agent be such that its contribution can aggregate toward coherent, integrated behavior? That is undoubtedly task and situation dependent, but to date there has been no unexpected or superior behavior from such aggregations. One level of agent granularity is shown in SRI’s Open Agent Architecture depicted in Figure 5.

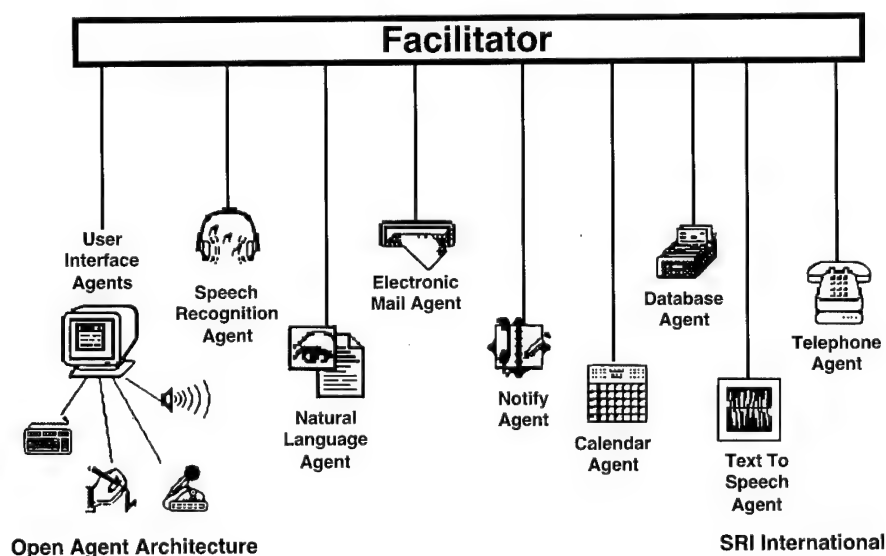


Figure 5 SRI's Open Agent Architecture

One area of development that could accelerate the use of agents in a wide range of settings is agent language; that is, a common agent language or environment containing user-agent, agent-agent, and agent-host (server) interaction. Several procedural, interpretive, and declarative (e.g., ACL/KIF/KQML) languages are already in use in new companies or in universities. Reducing agent functionality to a simple common form should, as in the reasons related to object-oriented software, simplify language and interactions and at the same time make trusted interaction more likely. Network-based operating systems that provide a homogeneous docking interface may also act to protect both host and agent. Common, public rules for agents that promote desirable inter-agent and host-agent behavior will be forthcoming. Scripting languages having just this purpose are now in use.

Regarding the underpinnings or catalysts to the growth of agents, several areas are worth noting. These will be expressed in human-like attributes even though the capabilities will be far from human. This illustrates the strong need for commonly accepted terms to describe agent attributes. Some agent properties:

- Veracity - Some availability of trust through tightly scripted interface language and cryptographic authentication.
- Competence - Some capacity for accurate observation and interpretation, limited by cost, to construct, maintain, and operate agents.
- Persistence - Recovery using some form of reasoning.
- Security and Safety - Some assurance through host constraints and agent veracity measures above. But two strategies will have to be avoided: 1) broad or intricate activities carried on by an agent in a host machine and 2) any true universal solution not supported by cryptography. [A worthy attribute: An agent's design is made such that its behavior from a host's perspective gets rewarded when it conforms exactly to its advertised purpose and punished, perhaps annihilated, when it does not.]
- Autonomy - In traveling agents, some "en route" decisions using reasoning. Autonomy is largely the province of the reasoning and learning agents; that is, being able to assess the situation and take alternative action and, in the latter case, remembering to avoid it next time.

Some applications of software agents are under development or test and illustrate the directions future design will take:

- Open Agent Architecture - Implements multimodal, distributed HCI (SRI)
- OASIS - An air traffic control system under test in Australia (AAII)
- FLiPSiDE - A "blackboard system" for agent interaction (Stanford)
- Telescript - A proprietary, commercial, general purpose scripting language (General Magic)
- SmalltalkAgents - A scripting language based on Smalltalk (Quasar)

- Tool Command Language (Tcl) - High-level, hypercard-like, machine-independent scripting language (public domain)
- Safe-Tcl - Secure version of Tcl (First Virtual).

5.5 Realizations in Software Agents Within 20 Years

The following projections are evolutionary, not revolutionary, and they begin with an assumption that the Internet will lead, directly or indirectly, to a global electronic information and commercial infrastructure. That infrastructure will be a consistent, universal, and pluralistic system. It will permit your personal or corporate computing environment, not necessarily local, to transparently represent you in a wide variety of transactions such as educational or learning systems, commerce (buying and selling), conferencing, scheduling, entertainment, mail, and much more. If such an infrastructure is not forthcoming, it will be for organizational and not technical reasons.

More about agent attributes:

- Veracity - Will be guaranteed through at least one of several methods: one-time authentication from a trusted third party plus checksumming; task execution or memory constraints in hosts; agent (of arbitrary complexity) surrounded by a simple shell written in script of constrained functionality; creation of very isolated environments in the host.
- Competence - Specified abilities to gain closure, accuracy.
- Autonomy - Specified circumscription but still wide latitude on agency and ability to negotiate or have volition. Agent-agent negotiations will occur under relatively simple, user-defined, and legal guidelines. Context will likely be buying and selling simple products rather than contracts in which both costs and benefits are more ambiguous.
- Delegation - The ability to receive abstract, human-language commands and carry them out in ways transparent to requester. It is not necessary to think of an agent as yours, as a single module, or as local. But roving, rogue, ownerless agents are an act of information warfare even if they have no malicious intent. Agents must always be responsible to some user!

5.6 Untethered Realizations in Software Agents

While twenty years seems an eternity in the computing world, there is one portrayal of the future of agents and HCI that has no time limit, only direction. This direction is from the present hands- and eyes-intensive machine toward a more amorphous system that might be called a delegatable assistant. These properties will emerge:

- Entire computers will become delegatable agents with natural language capabilities (See Human-Computer Interaction, Chapter 4.0)
- Trusted interactions will occur between users and hosts via their agents.

-
- Collaboration will occur among task or knowledge specific autonomous agents to achieve an integrated goal.

The notion of agents as delegatable software will occur as part of the evolutionary mainstream of computer development. Human-computer interactions will be done dominantly in human language terms much in the manner of requests or delegations. Keyboards will survive for text-intensive input.

6.0 Information Access technology

Introduction

Information Access Technology (IAT) provides the architectural foundation for Information Systems, that is computer-based systems which gather data and process data to provide information for decision-makers. Information systems also invoke tools to develop and assess alternate courses-of-action. These tools will operate over the global communication networks now emerging, and use a diversity of computer hardware and software to achieve their results.

Many information-processing tasks can be defined independent of specific source databases and of specific recipient applications. In the IAT architecture we assign such sharable services to an active middleware layer, and define domain-specific *mediator* modules to populate this layer. Mediating services must be of value to the customers, so that their applications will access mediators rather than sources directly. Several types of value can be considered: improvement in access and coverage, improvement of content, and delegation of maintenance. Criteria for mediating modules are: ownership by party who assumes responsibility for the results of the services, domain-specificity to delimit the scope of such a responsibility, and conformance with interface standards that enable interoperation.

Applications that benefit from mediation include planning and other decision-making tasks that require information from diverse resources (e.g., databases, reference systems, data obtained from sensors, and analyses projecting trends into the future). The sources are often autonomous, some of them are public. They will be heterogeneous. The heterogeneities include: representation, scope, level of abstraction, and context.

Mediated results are intended to be composable by higher level applications, which have to solve problems involving multiple subtasks. Hence mediators need a machine-friendly interface to support those applications. This interface must provide good communication, while encapsulating the domain-specific tasks, so that the complexity of the composed system is not much greater than that of the individual subtasks. However, questions of effectiveness and efficiency do arise, and must be dealt with by exploiting the processing and storage capabilities of modern hardware.

The corresponding architecture is a generalization of a server-client model. The partitioning enhances maintainability: the applications software staff can concentrate on functional improvements, the data resource managers on operational issues, and the technical maintenance are concentrated in the mediator layer. The concept is network oriented, and it is hoped that mediating services will be provided over the network by domain specialists. On networks, mediating services can be performed by independent entrepreneurs.

6.1 Information Access Technology

Information Access Technology (IAT) provides the framework and foundation services for accessing and shipping data of potential relevance to an Air Force decision maker. Central in this task are the conversion of data from the many heterogeneous data resources available for logistics, plans, local status, and intelligence into an integrated information base that provides situation awareness. In addition to transmitting data IAT needs to support ongoing access to software tools, which carry out the tasks. Some of the tools fall within the IAT domain, such as

rule-based interpreters which describe the needed data reduction for a particular application. Others are from their own support domains, such as image processing software for fusing intelligence data, which is invoked when needed for an IAT task.

As information systems are increasing in scope, they depend on many diverse, heterogeneous resources. These resources are typically developed and maintained separate from the majority of the applications that use their results. The applications that motivate the establishing of a database tend to be operational: inventory control, payroll, production control and the like. These data soon become important to support high-level applications: planning and decision-making. Decision support applications are typically designed subsequently and independently. Planning support is synchronized with objectives that change rapidly, and has to rely on existing sources, since it is rare that sufficient time is available to build planning systems and their data collection from scratch. Other sources of information are data systems such as digital libraries, geographic information systems, and simulations.

Transform Data to Information




Figure 6 Transforming Data to Information

Dealing with many, diverse, and heterogeneous sources overwhelm high-level applications with excessive emphasis on application-irrelevant but crucial details.

Since integrating many streams of data further increases the volume of data that can be delivered, a critical and implied task of IAT services is the reduction of this excessive volume to a manageable amount of essential information.

It is crucial to avoid information overload at the decision-maker by presenting only relevant information at the appropriate level of abstraction. That level will differ for an airman on the maintenance line, a pilot on a mission, or the commander of a joint task force. Information

overload also has to be avoided when lower bandwidth links are to be used, often in the last 100 miles to the warfighter.

6.2 Mediators

Mediators are the modules which provide intermediary services in an IAT architecture, linking data resources and application programs. Their function is to provide integrated information, without the need to integrate the data resources.

Specifically, the tasks required to carry out these functions are comprised of:

1. Accessing and retrieving relevant data from multiple heterogeneous resources
2. Abstracting and transforming retrieved data so that they can be integrated
3. Integrating the homogenized data according to matching keys
4. Reducing the integrated data by abstraction to increase the information density in the result to be transmitted

These are shown in Figure 6.

Advanced IAT services must also access simulation software, so that the decision maker can also develop and compare multiple courses-of-action. Overall, this processing adds value by converting data to information.

Definition: A mediator is a software module that exploits encoded knowledge about certain sets or subsets of data to create information for a higher layer of applications. It should be small and simple, so that it can be maintained by one expert, or at most, a small and coherent group of experts.

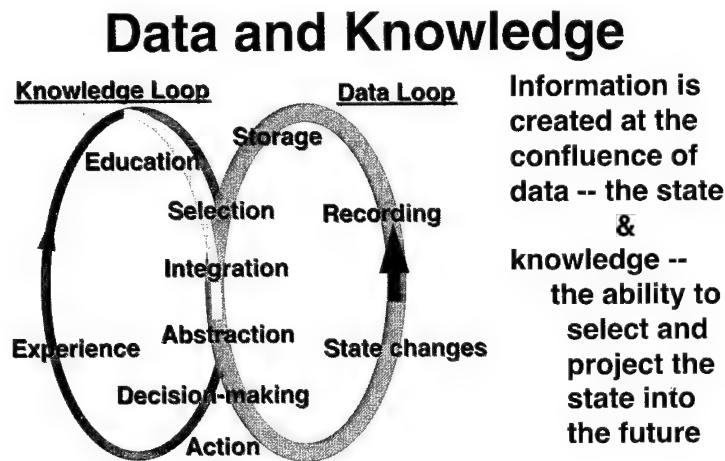


Figure 7 Interaction of Data and Knowledge

Access is a necessary prerequisite to deal with distributed information. Well-known examples of approaches to improve access are knowbots, facilitators, a variety of webcrawlers, and multi-database systems.

6.3 Module Types

We recognize two types of IAT modules: *mediators and facilitators*; IAT systems can be composed out of one or both of these. A mediator provides data processing functions, and is managed by a domain specialist who determines the functions and supports the needs of the recipient. The program in a mediator is typically an expert system, with knowledge about: the data sources; their significance; data transformations appropriate to the domain; and the information structure preferred by the recipient. A facilitator is a module which searches the network for sources of data, and delivers newly discovered data to a mediator and its maintainer, or directly to an end-user recipient.

In mediation more effort is devoted to processing the results of the retrievals (task types 2-4) than to access. Because of the processing effort, and the expectation that mediation services can support multiple applications, mediation is done in distinct modules in the communication networks, as sketched in Figure 8 (a logical progression of technology).

Mediated IAT technology depends on having adequate and secure communication networks, and builds on the technological infrastructure currently being established. Mediation technology augments the communication network with software modules, that typically will execute at intermediate nodes on the network, placed somewhere between the data sources and the computers used by the recipients of the information.

6.4 Status of This Technology

Mediation is an architecture which extends the client-server model and provides scalability. The sketches in Figure 8 show the development from:

- (a) Client-server applications in a setting that assumed compatibility (dealing in a single domain, often dominated by local data) to
- (b) The addition of wrappers to deal with legacy data and an integrator to combine resources, to

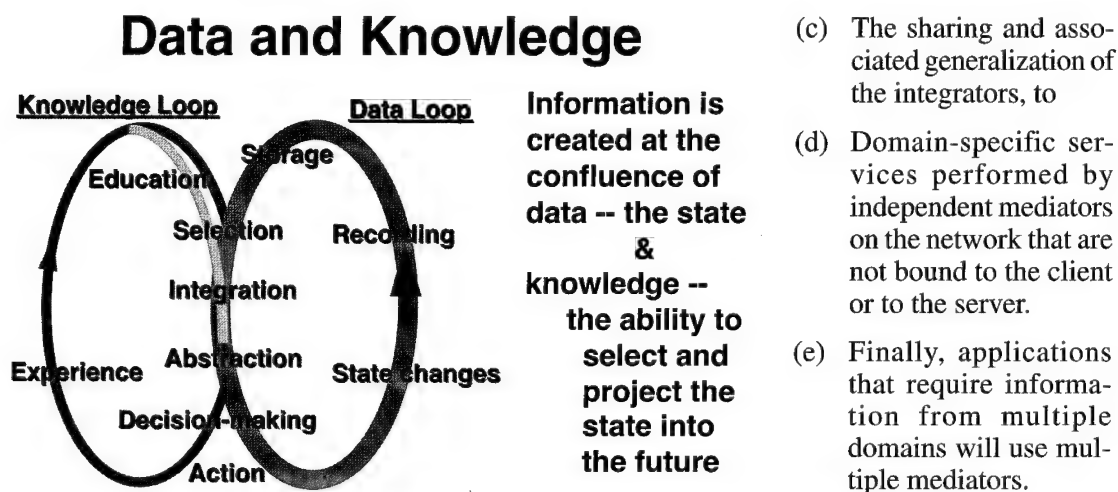


Figure 8 Network services from Client-Server to Mediation Layer

Mediators are being built now by doing careful domain knowledge acquisition and handcrafting the required code. Building mediators by hand is essential to validate the concept and establish the needed standards for the interfaces. Since the mediator approach is conceptually comprised of three layers, as shown in Figure 9, there will be two major interfaces:

- a: Base resources to mediation
- b: Mediators to applications

In practice there will also be intermediate interfaces, since it is likely that within the layer a number of sublayers exist as well. For these the technologies of type b. are also appropriate.

Functional Layer

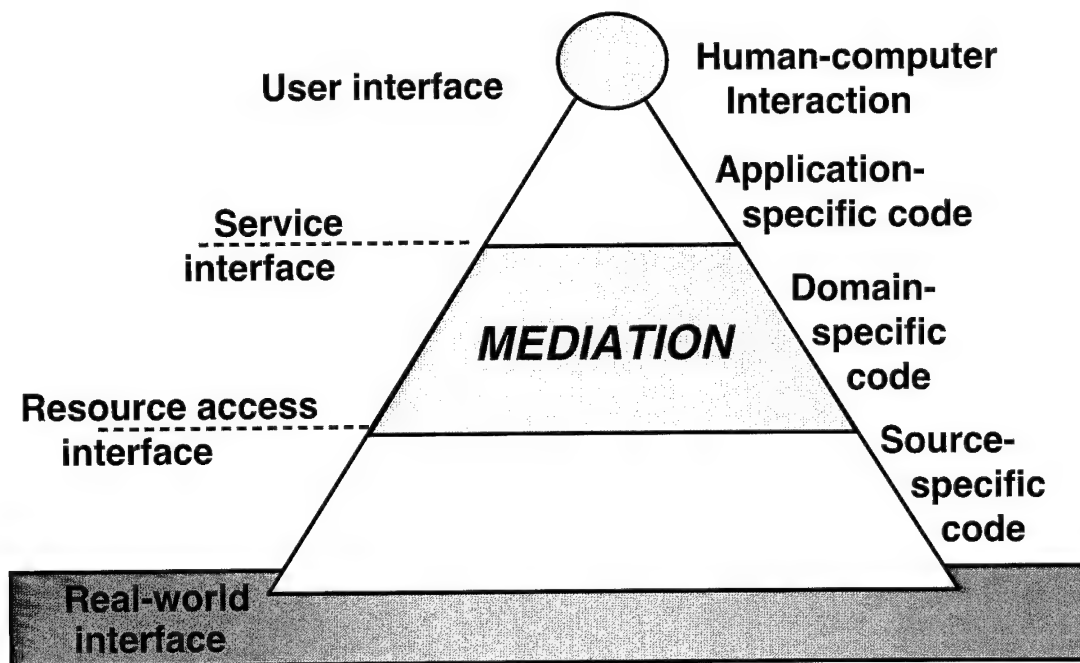


Figure 9 The Mediation Layer

6.4.1 Interfaces

For the base interface the many tools that are becoming available to serve the two-layer server-client model are appropriate. These include:

- a. Distributed and augmented SQL
- b. The interface tools for object-oriented access, such as the object fetching protocols of COBRA

At the application layer the interfaces need greater capabilities. Here agent languages are appropriate (See Intelligent Agents, Chapter 5). The sender and the receiver must agree on the chosen representation and on the vocabulary and its structure, i.e., the so-called ontology.

6.4.2 The User Interface

Mediation is simplified by delegating the complexities of the customer interface to the application program. Code to deal with the variety of graphical user interfaces (GUI), interface devices (screens, pop-up, roll-down, scroll, cut-and-paste, drag-and-drop, speech, etc.) often occupies more than 70% of application programs. Mediators and the invoking applications only need a machine-friendly interface, as represented by a KQML application program interface (API) block.

Trading internal code for an API recapitulates an earlier paradigm shift. We recall that in the 1960s the programming of file operations took a major amount of effort and competence. The acceptance of the database paradigm removed that code from applications and delegated it to database management systems (DBMS). The DBMS provides all the really difficult and specialized code associated with managing files, as backup, recovery, integrity, and internal consistency management.

Mediating modules can either provide information according to a pre-planned schedule, according to triggers that indicate that a significant change has occurred, or in response to a request. Principle functions that can be incorporated in a mediator include:

- a. Adaptation to the location and format of a data resources
- b. Selection of relevant data, aggregation of data into symbolic form to provide an effective abstraction for intelligent analysis
- c. Fusion of symbolic information according to knowledge-based rules
- d. A transformation to present the information to the recipients' computers for final integration of information from multiple domains and display of actionable items.

6.5 Services Provided by Intelligent Mediators

IAT systems provide services to information consumers. The nodes in the network that provide such services are the mediators. Value-added services in a mediator include combinations of:

1. Determination of likely resources using information extracted earlier
2. Invocation of wrappers that deal with legacy sources
3. Selection of relevant source material
4. Optimization of access strategies to provide small response times or low cost
5. Imposition of security filters to guard private data
6. Resolution of domain terminology and ontology differences
7. Resolution of scope mismatches
8. Interpolation or extrapolation to match differences in temporal data
9. Reduction of historical data to limited snapshots
10. Abstraction to bring material to matching levels of granularity for integration

11. Integration of material from diverse source domains based on join keys
12. Omission of replicated information
13. Assessment of quality of material from diverse sources
14. Pruning of data ranked low in quality or relevance
15. Omission of information already known according to the user model
16. Statistical summarization into higher level categories as defined in the user model
17. Reporting exceptions from expected values or trends
18. Triggering of actions due to exceptions from expected values or trends
19. Resolving temporal asynchrony in source data by referring to past, cached data
20. Completing current state information from past data by invoking simulation programs
21. Projecting status data into the future using simulations for Course-of-Action planning
22. Transformation of material to make presentation effective for the customer
23. Adaptation to the bandwidth and media capabilities of the customer
24. Transmission of resulting information and meta-information to the customer application

The conceptual underpinning for automating the mediation process is based on knowledge acquisition to provide the added value. Tools from the field of artificial intelligence that are a major contributor in mediation.

6.6 Where does Mediator Technology Stand Today?

Projects that are using mediators today have included manufacturing systems, such as the design of a new fighter aircraft at Lockheed; and gimbals for antenna positioning on spacecraft at Lockheed Space Systems. Early applications have been in integrating information for flexible military command systems. The technology for these projects is supplied by a variety of vendors, ISX corporation serves as a contact point for these projects.

6.7 Future Mediators for Effective IAT Systems

Technologies must be developed for automation of the building of mediators. Since the technology is modular there is promise in that direction, but the current short commercial funding cycles do not allow the development of the technology, since individual mediators, as now employed in demos and pilot situations can be more rapidly hand-crafted, one-at-a-time.

In order to build a mediator, resources have to be identified. This task is a common task carried out on the networks to enhance browsing, but not with the intent of building persistent paths and reusable software. Knowbots and facilitators provide the required base concepts. Once suitable data resources are located on the net, they have to be placed into a user-sensible hierarchy, as sketched in Figure 10. To achieve the required functionality, tools have to be attached to the nodes that provide for fusion and summarization. Some current ARPA funding is

focused on that goal, but it has to be sufficiently long-term to provide the vendor community with credible science.

Mediator Design Principle

Transform Data into
Information
Match User Model:
Hierarchical
to Resource Model:
General network
(and maintain models)




Figure 10 Mediator Design Principle

Beyond the initial creation of mediating modules is the need to update them. Where mediators are built with rule-based declarative technologies an excellent opportunity exists for automated adaptation, using machine learning technology developed in the artificial intelligence domain (See Artificial Intelligence, Chapter 7).

6.7 A Partitionable Architecture for IAT

The amount of knowledge to perform all possible aspects of mediation is enormous. Modularization or partitioning is the divide-and-conquer tool used in science to deal with excess size and complexity. Here size and complexity go together, because as the number of concepts grows, their linkages, their distinctions based on context, and their subsumption relationships also grow, and that growth is more than linear.

Since the number of domains is large we cannot foresee a central solution to that problem, and hence will have to deal with subsets of the domains, and bring those subsets together, as and when needed.

6.6.1 Domain-Specific Partitioning

Just as we specialize people in organizations, so we partition knowledge by domains. We define a domain to be a subset of knowledge in the world that can be managed by a single person, or at least by a small coherent group, as indicated in Figure 11. Within a domain no compromises should be necessary to define terms, and no committee effort should be required. It is best if there is an organization willing to take responsibility for a domain (e.g., the Defense

Logistics Agency for services supplying materiel; the Automotive Engineering Society for problems with cars).

The knowledge in a mediator module has to be maintained, since resources and user needs change over time. It is important to associate a mediator with a domain expert, responsible for the maintenance of its content.

Having the knowledge modularized in a domain-specific mediator allows convenient replications, since the mediators are programs that can be replicated onto multiple nodes on the network. The location of a mediator can be determined to optimize bandwidth, enhance security and be effective in terms of assignments of domain experts. The functionality of a mediator should not be affected by its location.

Specialization makes maintenance feasible. An expert owner can handle a specialized domain, without having to consider the different constraints imposed by handling unrelated domains, say, religion versus carpentry when discussing miters. Differing domains may also be best served by different programming paradigms, say finance versus engineering. Other subsets of domains may use similar technology, but differ in the concepts and structure of their knowledge representations, say electronic versus civil engineering.

A mediator is Not Just Static Software

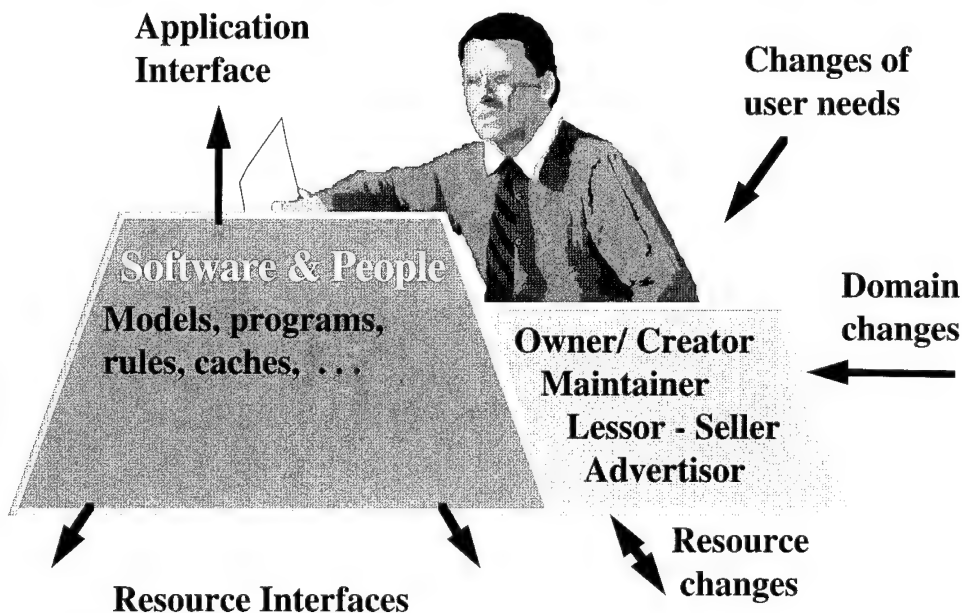


Figure 11 A Mediator has an Owner

Information from multiple domains still has to be integrated at the workstation of the the decision-maker, as indicated in Figure 12. The mediators provide functions that can be delegated to subordinate services, such as logistics and intelligence, but do not carry out the vital decision tasks of the commander in the field.

Integration at Two Levels

Application

- Informal, pragmatic
- User-control

Mediation

- Formal service
- Domain-Expert control




Figure 12 Integration at Two Levels

6.6.2 Distribution

Mediation is achieved by software. A mediator transforms data available on the network, to make it more suitable and relevant to the consumer. This software function can be carried out on the computer where the mediation was developed, or can be allocated to other computers on the network. Since software is easy to copy over the digital networks, mediators can be rapidly moved and replicated. Software is typically much smaller than the data it processes, so that it is easy to install at any location where compatible computers are available, as shown in Figure 13. Reallocation provides optimization for the IAT system, as well as an opportunity to tailor dataflow to the available bandwidth, a crucial concern once we leave the IAT backbones and proceed onto the difficult “last 100 miles” to the warfighter.

Allocation Flexibility

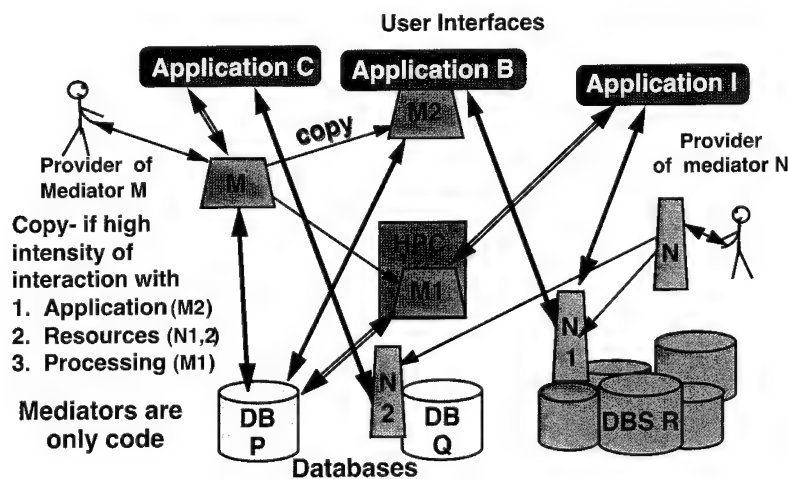


Figure 13 Allocation Flexibility

It is crucial that experience be gained with IAT concepts involving modern technology, so that rapid assembly of IAT systems will be possible when needed to assemble mission nets for new situations and rapidly assembled task forces.

Integration of simulation results will be defense-specific and must be integrated so that results of developing stand-alone simulations can become a part of ongoing tactical support systems.

If the processing algorithms used in mediation are costly, the software can be moved to a high performance processor on the network, again improving response time. Some of the features associated with a mediator architecture are sketched in Figure 14.

Features of Mediation

- Domain-specific partitioning for Creation and Maintenance
- Network-basing for easy Reconfiguration
- Caching to deal with Asynchronicity
- Replication for Performance

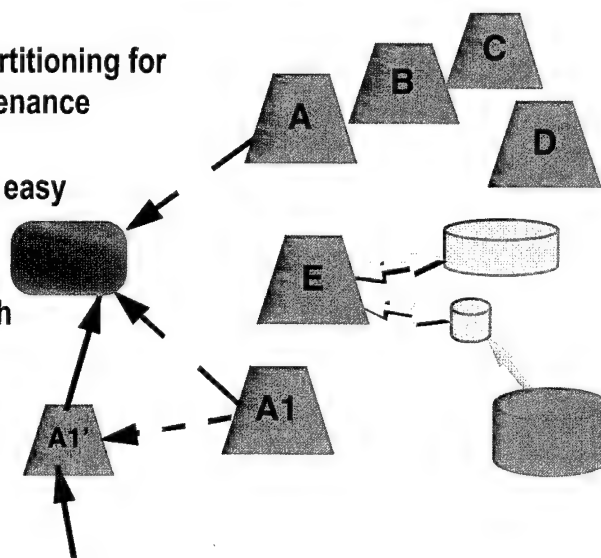


Figure 14 Some Features of Mediators

If demand for a particular type of mediator is strong, replicates can be distributed to additional sites in the IAT network. Since now the communication links will shorten, response time is enhanced as well. It is also possible to store acquired data or intermediate results to enhance the performance of mediators. For instance, detailed terrain maps are best kept on a node close to the customer, so that only changes have to be transmitted when they occur.

6.6.3 Standards

Since we assume that most mediation services are performed in autonomous computing nodes, the requirement for interoperation is the ability to communicate according to some standard conventions. Heterogeneity of computing platforms, operating systems, and message passing infrastructures is being overcome by concerted efforts in many communities. For mediation we must also consider terminology and representation conventions.

A number of projects are now using the mediator concept, and some standards (KQML is an example) are being promulgated for the interfaces that are required. A series of prototype efforts supports the US Air Force concept of an Integrated Weapons Systems Database. A phase is sketched in Figure 15. These prototypes allow for the gathering of experience that can validate the proposed standards. Investment will be needed to provide commonality for that effort, since there is a tendency for independent suppliers to compete in terms of leadership in standards as well, disabling the objective of vendor independence for the purchasers.

F-22 IWSDDB Phase 6

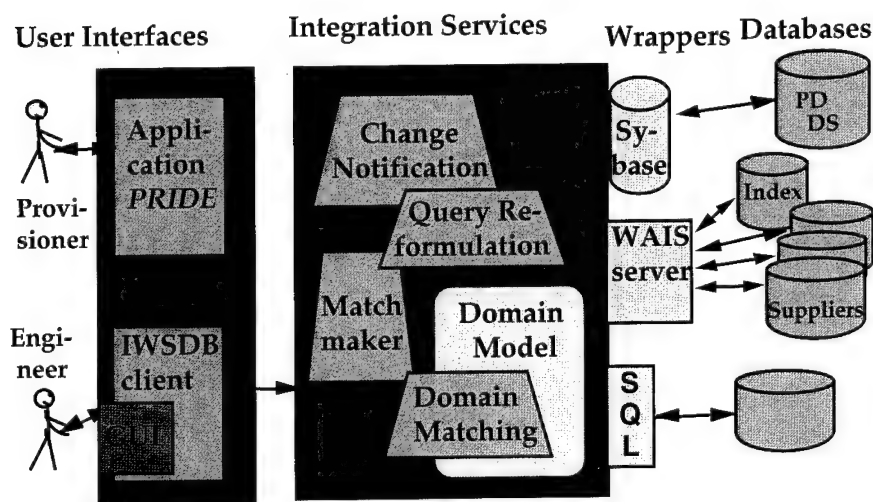


Figure 15 F-22 Integrated Weapons Systems Database

6.6.4 Maintenance

Mediation adds value to the data by applying the expert knowledge of the expert who has created the mediator. Mediators should also be maintained by those experts, so that the mediators remain effective in a constantly changing world. As soon as an improved mediator is developed it can be advertised over the network, both to existing subscribers as well as to potential new clients. A poorly maintained mediator will lose value over time, and be a candidate for replacement by a competitor.

6.7 Conclusion

Information technology is serving us well in specific domains, although we have remained dependent on specialist model designers and programmers for our implementations. Object technology has lessened our dependence on specialists by being able to use an infrastructure that aggregates detail into meaningful units.

Mediation is a technology that is intended to scale systems so that sources from many domains can contribute services and information to the end-user applications.

7.0 Artificial Intelligence

Introduction

The Panel's projection of the future of the software technology called Artificial Intelligence (AI) was greatly helped by the timely appearance of a paper on the subject, prepared by a distinguished group of AI scientists, co-chaired by SAB and panel member Davis and organized by the American Association for Artificial Intelligence. The first order of business for the reader who wants to better understand the future of AI is to read this paper ("A Report to ARPA on Twenty-First Century Intelligent Systems", AI Magazine, Fall 1994). We quote liberally from this paper in the material below.

Advances in computers and telecommunications have made a vast quantity of data available to us, and given us computational power that puts the equivalents of mainframes on the desktop. However, raw information processing power alone, like brute strength, is useful but insufficient. To achieve their full impact, systems must have more than processing power—they must have intelligence. They need to be able to assimilate and utilize large bodies of information, to collaborate with people and to help them find new ways of working together effectively. The technology must become more responsive to human needs and styles of work, and must employ more natural means of communication.

To address the critical limitations of today's systems, we must understand the ways people reason about and interact with the world, and must develop methods for incorporating intelligence in computer systems. The concepts, techniques, and technology of the IT area called Artificial Intelligence offer a number of ways to discover what intelligence is—what one must know to be smart at a particular task—and a variety of computational techniques for embedding that intelligence in software.

Below we describe AI applications and underlying technology that will enable intelligent systems to meet Air Force needs in the next five to twenty years. We will refer to these applications as "high-impact application systems."

7.1 Intelligent Simulation Systems

Elsewhere in this volume (See Modeling and Simulation, Chapter 9) we offer a view of the future of modeling and simulation technology. A new generation of *intelligent* simulation capabilities will support the construction of programs that model complex situations, involving both complicated devices and significant numbers of intelligent simulated people. The simulated worlds that can be generated today have limited physical realism and severely lack realism in their simulations of people.

7.1.1 Artificial Simulation Actors

The systems of the future will differ in both scale and function from those that exist today. In the next generation or two of simulations, thousands of "actors" will play roles. It might be economical to use actual people for only a few of these roles; the rest could be simulated using AI techniques.

A key challenge is constructing realistic humanlike actors. In the future, these actors will be able to coordinate perception, planning, and action (discussed later), learn, understand and interact with their world, deal with other actors, and use natural language. Providing all or even a significant portion of this functionality is a challenging mission. However, useful agents can be constructed with only some of these capabilities—even in limited form.

7.1.2 Simulation in Engineering

Another quite different use for intelligent simulation will be for an advanced form of engineering design and evaluation. AI programs can be given the knowledge and reasoning power to use the physics and engineering principles underlying the design of artifacts. Such programs will speed up and make more accurate the formulation of design models that can be tested by conventional numerical simulation. An evaluation environment for new products, such as vehicles or airplanes, could use simulations of people to test the feasibility of a product's construction, use, and maintenance before it has been built. A new product design could be "used" by simulated people while it has only a virtual existence. Potential customers could try out the product in a simulation.

7.2 Intelligent Information Resources

Information-resource specialist systems will support effective use of the vast resources of the national information infrastructure. These systems will work with their users to determine users' information needs, navigate the information world to locate appropriate data sources—and appropriate people—from which to extract relevant information. They will adapt to changes in users' needs and abilities as well as changes in information resources. They will be able to communicate in human terms in order to assist those with limited computer training. These systems constitute an important class of *intelligent agents*, discussed in Chapter 5.

7.3 Intelligent Associate Systems

Software designed to act as an intelligent, long-term team member could help to design and to operate complex systems. An *intelligent associate system* can assist with design of a complex device (such as an airplane) or a large software system by helping to preserve knowledge about tasks, to record the reasons for decisions, and to retrieve information relevant to new problems. It could help at the operational level to improve diagnosis, failure detection and prevention, and system performance. Associate systems do not need to be experts themselves; rather, they could significantly boost capability and productivity by collaborating with human experts, assisting them by capturing and delivering organizational memory.

The Boeing 777 aircraft illustrates that some major advances in design technology have already taken place. New tools enabled designers to check spacing and clearance so accurately that a physical mock-up version of the plane was not needed. But these tools still had limitations. They did not incorporate, for example, vast volumes of design information. As a result, engineers had to manually consult printed documents. Other information, such as some of the compromises made in the design process, was never recorded, has now been lost, and will be sorely missed when the design is revised in the future (as all designs are).

7.3.1 Intelligent Help for Information Overload

Sensor and communication systems provide the warfighter with a wealth of data for decision making. In the future this wealth threatens to be overwhelming. The clearest uses of Intelligent Associates will be to assist individual users and teams to gather, cull, organize, and interpret data relevant to a situation. The Information Applications Panel discusses the future of Information Fusion. The AI technology to make this vision a reality will mature in the next 10-20 years.

A recent report by the Office of Science and Technology Policy noted that in the near future, every home and business could have an information appliance that combines the capabilities of telephone, television, newspaper, computer, and Internet services such as electronic mail. The translation to a warfighter's workstation is obvious.

7.3.2 Intelligent Help for Ease of Use and Communications

To realize this enormous potential, Intelligent Associates must be powerful, flexible, and easy to use. Users must be able to communicate in whatever way is most natural to them: typing or speaking, for example, in their native language rather than some artificially designed language. Associates will allow the use of diagrams and gestures, combining media and modalities in whatever mix is best for getting the message across. The commands that users issue will be general and often vague; nevertheless the Associate must accurately determine how to perform such commands. The information that a user needs will often not be stored at any one site; thus the Associate will need to be able to access multiple sites and recognize common information (see Chapter 6 on mediators). To actively and continuously seek out useful information, an Associate will need to learn which topics are of long- and short-term interest to each user.

7.3.3 Intelligent Help for Organizational Processes

The Associate will remember and recall the rationale of previous decisions, and, in times of crisis, explain the methods and reasoning previously used to handle that situation. Intelligent Associates will incorporate intelligent simulation and information resources systems as components.

For example, an Intelligent Associate for aircraft design will enhance collaboration by keeping communication flowing among the large, distributed design staff, the program managers, the customer, and the subcontractors. It will also assist in adapting existing design during modifications and subsequent generations; support concurrent simulations of an overall design whose components might be in various stages of completion; and capture design rationales (such as for wing design), making them readily available during the entire design lifetime and accessible for maintenance and repair.

7.3.4 Intelligent Help for Software Development

One critical area in which Intelligent Associates will assist is software development: keeping track of specifications, design proposals, and implementations for a software project throughout its life cycle; recording the design decisions of a constantly changing team; and being a repository of solutions and components for new projects. For the new architecture-based software development, interactive AI methods will be used to instantiate requirements and specifications

as a bridge to an automatic coding process, and bring many “special case exceptions” into the code (from a case library).

7.3.5 Intelligent Help for Finding Analogous Situations and Cases

The Intelligent Associate will use methods for reasoning by analogy. Analogy techniques could be used to look for existing specifications, components, or implementations that match some new requirement.

7.3.6 Intelligent Help for Managing Complexity

Intelligent Associates will assist with many of the problems that arise in using our ever-more-complex systems, including diagnosis, planning, and operational tasks. For example, they will add significant value to the operational control of air vehicles and weapons systems. During both normal operations and emergencies, the Associate will monitor information derived from sensors in the control arena or cockpit, providing guidance and advice based on previous experience to the warfighter.

7.4 The AI Technology Underlying High-Impact Applications

A common core of capabilities is needed to construct AI applications. These include:

- a. Abilities to reason about the task being performed with the knowledge that is appropriate to the task.
- b. To reason about the collaborative process and the knowledge and capabilities of other systems and people participating in an interaction.
- c. To communicate with users in human terms, producing and understanding combinations of spoken and written language, drawings, images, and gestures.
- d. To perceive the world.
- e. To coordinate perception, planning, and action.
- f. To learn from previous experience and adapt behavior accordingly.

Understanding these capabilities in humans and developing computational techniques to embody them in programs has been a central focus of AI research. A solid foundation has been developed in the large body of previous research. This work produced the technology that underlies the few thousand knowledge-based expert systems used by industry and the Armed Services, as well as many other applications in planning, learning, perception, and language processing.

7.4.1 Learning, Automatic Adaptation

Virtually all high impact application systems can be more powerful if they can learn from experience. For example, Intelligent Associates that can learn will be able to tailor their information retrieval process to a user’s needs without having to be told exactly what to do. They will instead generalize from previous interactions with the user. Learning skills will enable

an Intelligent Associate to deal with new types of problems, for example, drawing on its experience in the design of one type of UAV and applying it to the design of another.

Basic research has steadily advanced the fundamental technology of machine learning for more than two decades. A wide variety of learning methods—including decision-tree induction, neural networks, genetic algorithms, explanation-based learning, and case-based reasoning—have empirically demonstrated their utility on a broad array of real-world problems. There have been significant advances. These include:

- a. Goal-directed learning, in which programs make decisions about what, when, and how to learn.
- b. Practical methods for learning in the presence of a significant number of irrelevant features.
- c. The use of knowledge the system already has to improve the quality of learning.
- d. Use of machine-learning techniques for scientific discovery and other kinds of data mining.
- e. The integration of learning with planning, language processing, and perception-action.
- f. Active learning, in which programs design experiments and other information-gathering activities that supplement the analysis of presented data.

In the future, neural networks will develop to be powerful pattern classifiers and will be used as “front-ends” for symbolic reasoning programs. Genetic algorithm methods will slowly develop as relatively weak search methods to aid machine learning.

7.4.2 The Plan-Decide-Act-Monitor Cycle

Intelligent systems must be able to plan—to determine appropriate actions for their perceived situation, then execute them and monitor the results. Planning, in turn, requires advanced capabilities to represent and reason about time, action, perception and the mental states of other agents. To cope with realistic situations, systems must be able to deal with incomplete, uncertain, and rapidly changing information and must have mechanisms for allocating resources between thinking and acting.

7.4.2.1 Planning

Basic research in planning has provided a substantial base on which to develop intelligent planning capabilities. A variety of algorithms have been developed for constructing plans to satisfy a given set of goals. Learning techniques have been applied to reduce the time planners take to solve problems by enabling them to effectively apply previously derived solutions to new problems. Recently, a new class of planning systems was developed that combines perception, planning, and action and guarantees a response in bounded time. These “reactive planning systems” function in dynamic worlds to which they are connected by their perceptual system; they are more easily linked to traditional control mechanisms for the low-level operation of effectors.

Practical systems that have been crafted to take advantage of domain-specific constraints can automatically develop plans consisting of thousands of actions, both sequential and parallel, in domains such as logistics and battle planning. New capabilities will manage the trade-offs among acting, planning, and acquiring further information to reduce uncertainty.

Since the technology for planning from AI and from Operations Research is highly developed, we can forecast that over the next 10-20 years most human decision making involving complex sequences of actions and parallel courses of actions will be assisted by (at least) semi-automatic computer planning

7.4.2.2 Perception and Language

The ability of computer systems to perceive and communicate has evolved dramatically over the past decade. Large-vocabulary, discrete-phrase speech recognition is commercially available; several laboratories have developed speaker-independent real-time continuous speech recognition systems for tasks requiring several thousand word vocabularies. These systems will be rapidly commercialized. We can expect to see in the coming decade, highly effective continuous speech understanding systems, with tens of thousands of words, with error rates not exceeding 1-2%. Such systems will ride the wave of increasing computer power available cheaply.

The systems will complement advanced natural language-processing techniques, which now support automated clipping services for categorizing newspaper stories, as well as partially automated translation of technical manuals into foreign languages. Such applications for automatic or partially-automatic natural language understanding will become commonplace in 10-20 years.

Growth in the coverage (scope) of natural language understanding, and its reliability, will track the growth of the large knowledge base described later, with high performance capability in specialty areas occurring in five years. However, for unconstrained (but unsophisticated) human discourse, the time frame is more like 20 years.

7.4.2.3 Perception

Significant technical progress will enable real-time perception with acceptable accuracy. The methods being investigated in the perception community include using more sources of information, and designing automatic training methods that work alone or in combination with hand-crafted rules and models. Image understanding techniques are being developed to interpret multiple views of the same scene or event, for example, in a video of an object in motion.

Symbolic rules and models will be augmented by methods that learn automatically from data the likelihood that a rule or model component will be applicable in a given situation. These techniques, which take advantage of informative statistical patterns that humans cannot reliably detect, will improve the robustness of the interpretation process and decrease the time necessary to adapt a perceptual system to a new domain.

Finally, central to all the perceptual modalities is how to coordinate symbolic methods with nonsymbolic ones (for example, stochastic methods or neural networks). Research has reached the stage where significant advances in the technology will occur that will allow the combining of the best features of both approaches.

7.4.2.4 Human-Computer Communication in Multiple Modalities

Communication among people is marked by its flexibility, from the casual nod of a passerby conveying a greeting, to a professor's math lecture with its complex interaction of lecturing, drawing diagrams on a chalkboard, and answering questions. People use a number of different media to communicate, including spoken, signed, and written language; gestures; sounds; drawings, diagrams, and maps. The high-impact application systems must also be able to understand the full range of communication media.

For example, an Intelligent Associate helping a warfighter might provide information using a combination of maps, diagrams, text, and spoken descriptions. These various media will be combined so that information is communicated in the manner most appropriate to the particular user and task at hand.

Interpretation and synthesis processes in individual modalities are subject to a certain degree of error; even humans misunderstand each other. The joint use of multiple modalities permits one modality to compensate for interpretation errors of another.

Broad-band human-computer interaction has the potential for large payoffs. Techniques for fusing multimodal input will serve as the basis for simpler interfaces that allow the user to combine pictures, speech, mouse, and keyboard input, using each where it is most convenient. It is either ironic or amusing that we will be using one of our most powerful (software) technologies to simplify the use of our complex IT systems.

7.4.2.5 Finding Something by Its Content

The Internet is already populated with enormous amounts of multimodal information, from pages containing images, text, and graphics to video with sound track. This wealth of information will grow ever more extensive when the NII and the DII (Defense Information Infrastructure) become realities. Intelligent systems will provide access to a wide variety of information, including visual and audio data, in addition to commonplace structured databases.

Any access to these materials beyond the simple keyword and hypertext browsers now available will require automatic indexing schemes that work across multiple modalities and will require capabilities for content-based retrieval. Recognition of moving images of objects in video material, a substantial benefit to analysis, will be required in the future and will be available.

7.4.2.6 The Power to Reason

In any realistic problem, reasoning must be done under less than perfect conditions. Intelligent information systems must deal with data that is imprecise, incomplete, uncertain, and time varying. They must be able to manage with domain knowledge that is incomplete, and they must do so as they meet pressing real-time performance requirements. Finding a solution that is guaranteed to be optimal—under any reasonable interpretation of optimal—can be shown to be computationally intractable, i.e., cannot be done efficiently no matter how much faster we make our computers. Consequently, we must develop fast methods for plausible reasoning that can be shown to lead to good—if not necessarily optimal—solutions.

Sensor data providing the system with recent information may be imprecise and, from time to time, unreliable because of sensor failures, drifts, or extreme operating conditions. This

incomplete and vague data must be reconciled, integrated with available statistical information, and analyzed to identify trends and situations that require corrective actions. Decisions must be made quickly and in a way that can be justified to the end-user.

AI research to date has partially addressed these issues by developing many specialized reasoning techniques, including:

- a. “Anytime” reasoning, techniques for enabling a system to reach the best possible conclusion within the time available.
- b. Nonmonotonic reasoning, techniques for leaping to conclusion based on partial information in a justifiable way that allows conclusions to be withdrawn if necessary as new information comes in.
- c. Case-based reasoning, techniques for using previously acquired solutions to old problems as the basis for new solutions to new problems.
- d. Bayesian networks, techniques for using causal and probabilistic information efficiently.

Techniques based on probability or inexactitude fall into a class that the Japanese now call “soft” (as opposed to the purely logical “hard”). Also in this class are “plausible” reasoning methods based on heuristic knowledge, and reasoning based on “fuzzy” set definitions with membership ranges. These methods will probably be the most widely used methods of reasoning in 10-20 years (vs the “hard” methods). For the larger world of computer applications, these “soft” reasoning methods may become more important than calculation (numeric computation).

7.4.2.7 Representation

A variety of representations that capture information at multiple levels of abstraction and in different degrees of detail will allow programs to reason effectively about complex systems. For instance, the most abstract level will represent the core conceptualization, providing information about the way an artifact accomplishes its goals. Programs will be able to reason quickly, but only imprecisely with representations at this level. More specific representations will encode more details and enable more precise reasoning, but with greater computational cost and increased difficulty in interpretation.

7.5 The Scaling Up to Large AI Systems

The most important limit to the intelligence of current AI systems is their narrow scope and shallow depth. These systems have not been given (by us) enough knowledge to assist us in the great variety of our tasks. Nor have the systems yet learned large bodies of knowledge by machine learning methods. In the next decade, major advances will be made in producing an international distributed knowledge base of the “widely-shared knowledge” of our society and our science. (Think of this as the “Knowledge Web” in analogy to the World Wide Web.) The knowledge web will be the backdrop to all AI systems. Networks will be the medium by which this knowledge is accumulated and distributed. The commonly held view that “you have to tell a program everything” in order for it to perform properly will become obsolete. The “widely

shared knowledge" base coupled to reasoning programs will supply necessary but missing detail.

Brittleness has been a perennial problem with the thousands of expert systems constructed to date. They are good at their task but their performance falls off drastically as they move away from that task. Human expertise is far more flexible; it rests on a large stock of the previously mentioned widely shared knowledge about the world. The large knowledge base will solve the brittleness problem for most AI applications by providing this "fall-back" knowledge.

In the 10-20 year future, we believe the knowledge web will contain tens of millions of objects, rules, and logic formulas (perhaps hundreds of millions). A variety of reasoning packages will be available to plug-and-play with the knowledge web; and will be easily customizable to the users' needs.

The combination of natural language understanding methods, machine learning methods, and the knowledge web may give us a powerful surprise. Knowledge-based learning methods may grow the knowledge web in a "bootstrapping" way by understanding natural language text (similar to how we come to know things by reading about them). For specific domains of application, such a combination will reduce the domain knowledge acquisition time by factors of ten to one hundred.

7.6 AI Technology Points of Interest

The technical requirements of significant Air Force applications are considerably broader than AI technology alone can provide. However, AI capabilities will be key to making Air Force systems intelligent, adaptable, far more accessible to the relative unskilled user, and, thus, dramatically more effective.

Providing AI capabilities is not an all-or-nothing proposition. Although the development of systems with very sophisticated capabilities will require long-term effort, in each category of application more restricted but still usefully intelligent systems can and will be developed.

7.7 Interdisciplinary R&D

High-impact AI applications require coordinated efforts of research and development across the several areas of computer science. Building these systems will require combining AI methods with non-AI approaches and embedding AI technology within larger systems. In addition, many of the fundamental scientific challenges require collaborative, interdisciplinary efforts in the cognitive sciences and engineering.

7.8 Summary

Artificial Intelligence technology will provide the foundation for systems that can search large bodies of data for relevant information; help users to evaluate the effects of complex courses of action; and work with users to develop, share, and effectively use knowledge about complex systems and processes. AI will make it possible to build a wide range of application systems that assist decision makers in adapting and reacting appropriately to rapidly changing world situations.

8.0 Computer-Aided Planning

Introduction

A number of research efforts currently underway will markedly improve the state of the art and practice in computer-aided planning over the next several decades. Given the scale and significance of planning operations in the Air Force, even evolutionary-scale improvements can have enormous significance. We see four threads that will have important impact on planning.

8.1 Constraint-Based Techniques Will Vastly Speed Up Planning

Recent laboratory experiments have demonstrated significant speed improvements with a technique called constraint-based scheduling. This approach speeds the search for a schedule by incorporating into the search process knowledge about the constraints the schedule must meet (e.g., ordering of events, intervals between events). Recent lab experiments have created automated schedulers by combining a powerful general purpose scheduler with constraints from a specific task; the result is a single, new program capable of very high speed generation of schedules that meet the specified constraints.

We believe that work of this type will continue to provide advances, resulting in an additional two orders of magnitude improvement in the speed of schedule generation, permitting substantially more schedules to be considered during a planning operation, offering more and better options to command personnel.

8.2 Rationale Capture Will Enable Powerful Planning Systems

Detailed operational plans are often similar enough to one another to produce in planners a sense of déjà vu and frustration (“Didn’t we just do something like this last month?”). Yet the plans are typically different enough to prevent straightforward reuse, resulting in a major (re)planning effort that repeats some substantial part of what may have been done recently.

One development that will help solve this problem is the capture of plan rationales, i.e., plans that indicate not only what to do, but what they were trying to achieve—their objectives—and why the actions given were selected as a way of achieving the goal (e.g., arguments for and against a particular action as a way of accomplishing an objective, underlying assumptions (such as troop strength) and estimates (e.g., likelihood of response)). Plans annotated in this fashion will have two important properties. First, they can be retrieved based on what they were trying to do, hence previous plans can easily be selected for review when a similar objective is encountered in the future. Second, where an off-the-shelf plan differs in its objectives from the current set of objectives, the off-the-shelf plan can easily be modified. Because actions are linked to objectives, such plans will have a kind of spreadsheet-like character to them: if objectives are deleted, the corresponding actions are easily identified, and if those actions now serve no other objectives, they can themselves be deleted. Because actions are annotated with explicit assumptions and arguments, where those assumptions or arguments differ, appropriate modifications to the plan are easily identified. An early version of this capability is the Air Campaign Planning Tool (ACPT) of the Air Force Pentagon Checkmate Office.

Such power will require development of sophisticated plan rationale description languages and advances in the technology developed to date to support only-line argumentation. While it

is relatively easy to imagine plans annotated with the sort of information described above, there is a formidable task involved in developing a language that a computer program can use to express, track, and reason about plans with nearly the same facility that is possible with spreadsheets containing numbers.

We suggest that within five years such systems will enable semi-automated planning and plan modification for applications such as transportation and force planning, using case-based reasoning, within the scope of one course of action. Within ten years we predict 90% automated plan selection and modification, over the scope of multiple courses of action.

8.3 Advances in Decision Theory Will Enable Adaptive Planning in Uncertain Environments

Advances in decision theory and reasoning about uncertainty will enable execution-time plan modification even in the highly uncertain environments characteristic of conflict. Numerous studies and extensive experience have shown that humans routinely err when faced with the need to reason with uncertain information. Advances in automated reasoning in this area have included developments such as belief nets, which provide an effective mechanism for expressing likelihood's and interdependencies of events. Some work has also been done on adaptive planning, i.e., the interlacing of planning and execution that is necessary when (as is typical in military environments) plans must be adapted to changing conditions.

Within ten years the combination of these two technologies will provide the routine ability to do adaptive planning in the face of uncertainty. In this world, plans will be formulated to take into account probabilities of various hypotheses, selecting the best option given the current best guesses, and will be appropriately modified in real time as additional information comes in or unanticipated events occur.

8.4 Multi-Agent Planning Will Permit Large-Scale Continuous Planning

A variety of increasingly powerful multi-agent planners will permit the integration of concurrent multiple viewpoints and, eventually, permit planning to be done continuously by multiple teams of people and software agents working together.

In the next five years we will see coordinated, multi-agent planning, by which we mean multiple programs and people working on a planning problem at the same time, "checking in" with each other as appropriate to coordinate their efforts. This will result in part from progress in computer-supported cooperative work (CSCW), but will build in important ways on AI work to automate planning (CSCW is aimed primarily at facilitating human work). One important payoff here will be the ability to automate planning that takes into account multiple perspectives (e.g., warfighter, supply, doctrine, etc.).

Five years beyond that we see the advent of *concurrent* multi-agent planning, by which we mean agents (both human and software) that are planning while in continuous contact with each other's efforts and results. The result will be a considerable speed-up in multiple perspective planning.

As these basic advances are made, impact of this technology area will depend significantly on advances in modeling and simulation: the effectiveness of any planner (whether human or automated) depends to considerable degree on the level of detail available for describing the resources being scheduled (e.g., capabilities of various transport vehicles, size, weight, range of warfighting equipment, etc.). As a consequence, we believe that coordinated parallel effort should be put into developing appropriately detailed and powerful models of Air Force equipment, so that these advanced planning techniques have the relevant raw material to work with.

9.0 Modeling And Simulation

9.1 Computer Modeling and Simulation Technology

Computer modeling and simulation provides parallel-world capabilities for the full spectrum of Air Force activities, from research and development through analysis, acquisition, test, evaluation, production and logistics to education, training, and operations.

Computer modeling and simulation has been evolving its broad range of Air Force capabilities since the first use of computers for ballistics applications in the 1940's. As a result, the current Air Force inventory of independently-developed computer models and simulations has formidable interoperability and compatibility problems. These have been overcome in some limited domains (e.g., flight dynamics models) and in some medium-scale training-oriented distributed interactive simulations such as Simnet. They have also been overcome in point-solution demonstrations and exercises linking wide varieties of models, simulations, real equipment, and operators. However, significant broad, regular-use operational issues of simulation such as interoperability, rapid configuration, and verification, validation, and accreditation (VV&A) are just beginning to be addressed.

By 2005, with a significant level of Air Force effort, basic large-scale interoperability support (consistent data definitions and fixed interaction protocols) will be available for critical-mass configurations. By 2015, this will extend to include dynamic interaction and interoperability agents. By 2005, simple user languages enabling rapid composition of models and simulations will be available. By 2015, these will extend to support automatic configuration of models and simulations to address given decisions (e.g., choice of airlift capabilities).

By 2000, the Defense Simulation Internet will provide broadband support of point-solution "Louisiana Maneuvers" scale simulation. By 2010, it will be operationally robust and able to support regular exercises at this scale. Currently, VV&A technology consists of basic test suites and simple assertion checking (e.g., of conservation of energy, resources, etc.). By 2005, with Air Force investment, this can expand to simple mission-domain model checking and built-in-tests. By 2015, this VV&A technology can expand to domain model checking using automated agents and dynamic built-in-test, achieving much higher levels of credibility.

By 2015, the resulting modeling and simulation capabilities will enable combat operations options to be credibly simulated before and during combat, greatly increasing combat effectiveness. The same capabilities will enable continuous two-sided exercise of information warfare capabilities, honing Air Force pre-eminence in this critical area. With appropriate attention to acquisition restructuring (e.g., virtual competition groundrules), it will enable virtual system acquisition, or flexible migration from virtual to actual combat systems, with complementary closed-loop combat system exercise and improvement across the system's life cycle.

9.2 Computer Modeling and Simulation (M&S) Rationale

In a warfare situation, the force which has the opportunity to explore the most options prior to combat has a significant advantage. The phenomenal growth of computer, communications, and software technology has enabled the Air Force to achieve such advantages via improved computer M&S.

Computer M&S enables the combat commander, mobility commander, or acquisition manager to open up numerous parallel “virtual worlds” in which many possible scenarios can be played out, and their likely outcomes assessed prior to commitment to a course of action. Current computer, communications, and software technology have enabled these virtual worlds to become sufficiently realistic to support credible virtual combat involving hundreds of human operators. This has enabled these distributed interactive simulations to serve as extremely cost-effective exercise areas for combat training, as well as for combat systems evaluation.

Computer, communications, and software technology advances described in the other chapters of this volume will continue to magnify the power provided to Air Force personnel by M&S capabilities. Not only will individual M&S’s become faster and more realistic, but improving M&S interoperability technology will make composition of multiple M&S’s a routine rather than a special-effort activity. Downstream, even greater advantages will result from embedding M&S’s into a closed-loop learning and improvement process for both the M&S’s and the systems they are emulating. This ability will be particularly important for information warfare, strengthening both information defense and offense capabilities via “information wargaming.”

M&S capabilities increasingly empower all sectors of the Air Force. Computational fluid dynamics models, materials phenomenology models, and knowledge models empower Air Force research and development. Combat M&S’s, weapons effects models, and cost models enable better force structure management, and better requirements determination for new Air Force combat systems. Engineering models, planning and scheduling models, and manufacturing models enable more rapid, predictable, and controllable system acquisition and upgrade. Test and evaluation is improved via hybrid human-in-the-loop and hardware-in-the-loop simulations, simulator/stimulator testbeds, and post-test analysis capabilities. Education, training, and operations are improved via human-in-the-loop distributed interactive simulation, re-creation of historic battles, and assessment of simulated course of action outcomes.

9.3 Computer M&S Definitions

We follow the definitions in DoDD 5000.59, “DoD Modeling and Simulation Management.” It defines a model as a physical, mathematical, or otherwise logical representation of a system, entity, phenomenon, or process. It defines simulation as twofold: a method for implementing a model over time; and a technique for testing, analyzing, or training in which real-world systems are used, or where real-world and conceptual systems are reproduced by a model.

As another perspective, the 1993 “Report of the DSB Task force on Simulation, Readiness, and Prototyping” says, “Everything is simulation except combat.”

This section of our panel report focuses on *computer* modeling and simulation. It thus excludes physical and mathematical models unless they are implemented as computer programs. It also excludes purely live simulations: exercises with live forces and real equipment in the field. It includes purely computational models such as computational fluid dynamics models, combat effectiveness models, logistics models, and cost models. It also includes hybrid computational and live simulations. These can be hardware-in-the-loop, as with sensor and missile test environment simulations; or human-in-the-loop, as with aircraft flight simulators and distributed interactive battle simulations such as Simnet.

For M&S applications, we use the hierarchical classification scheme from the excellent 1994 Defense Systems Management College (DSMC) guidebook, "Models and Simulations." Engineering-level M&S's address such issues as design, cost, manufacturing, and supportability. They provide measures of performance (MOP). Engagement-level M&S's address evaluations of system effectiveness against enemy systems. They provide measures of effectiveness (MOE) at the system-on-system level. Mission/Battle-level M&S's address evaluations of multiple platforms performing a specific mission. They provide MOE at the force-on-force level. Theater/Campaign-level M&S's address outcomes of joint/combined forces in a theater/campaign level conflict. They provide measures of value added at the highest levels of conflict, sometimes called measures of outcome. The DSMC M&S guidebook provides further information on the usual levels of detail, time spans, outputs, and uses of M&S at each level.

9.4 Computer M&S Constituent Technologies

Gigabit communications technology enables high-fidelity, large-scale distributed interactive simulations. Combined with flexible networking technology, it enables such simulations to dynamically evolve their combinations of real and simulated systems. By the year 2000, the Defense Simulation Internet will provide near-gigabit communication and networking support of point-solution "Louisiana Maneuvers" scale simulation. By 2010, it will be an operationally robust multi-gigabit system able to support regular exercises at this scale.

The massively parallel computing capabilities enabled by hardware and architecture technology enable the Air Force to better solve its "Grand Challenge" computational problems, such as signal and image processing, computational fluid dynamics, and weather prediction. Possible biocomputing breakthroughs will initially provide orders-of-magnitude improvements only for niche problems, but more general capabilities are also possible. Continued exponential growth of individual processor speeds and memories will enable massive mainframe M&S's to be exercised on individual desktops, with increasing fidelity and ability to explore multiple options.

Distributed software "middleware" capabilities will enable distributed simulations to be rapidly and reliably composed, including smooth adaptation to increasingly powerful commercial-off-the-shelf (COTS) software capabilities. The potential of massively parallel computing will be increasingly realized via improvements in parallel software: programming languages and compilers, operating systems and run-time services, and data management capabilities.

Increases in speed, flexibility, and scalability of Information Access Technology will strongly benefit large-scale M&S's, which are prodigious consumers and generators of information. Particular M&S functions which will benefit are multi-database interoperability and interactive analysis of M&S outcome data.

Software development technology will provide stronger software architecture support of families of M&S's, and the ability for M&S users to use very high level languages to compose and modify M&S's. In general, advanced software environments will enable faster, cheaper, and better M&S development and modification. In particular, technology will provide much-needed improvements in the design and testing of distributed and parallel software and systems.

Artificial intelligence technology will provide considerably stronger semiautomated force (SAFOR) capabilities for distributed interactive simulations. Speech and natural language technology will improve both human interfaces to M&S and the ability to analyze histories of human-in-the-loop simulations. Domain knowledge modeling technology will enable M&S's to be automatically or semiautomatically generated from statements of the decisions needing support (e.g., "I need to analyze the use of military vs. commercial aircraft for the following transportation scenarios. Configure the appropriate performance and cost models and provide me some comparison summaries on delivery time and cost").

Agent technology will facilitate M&S interoperability by use of agents to mediate inputs, outputs, control, and synchronization of multiple M&S's. Agents will also provide powerful information search capabilities for both configuring and operating M&S's. They will provide stronger assertion-checking capabilities for M&S verification, validation, and accreditation.

Assurance technology will also strengthen M&S verification, validation, and accreditation capabilities. In an information warfare context, it will provide stronger defenses against contamination of M&S software and results.

Improvements in human-computer interface and collaborative computing technology will dramatically enhance M&S capabilities for planning, development, testing, training, and operational support. Critical technology elements include high resolution displays, 3D displays, voice and speech recognition, speech generation, and tactile interfaces. For example, a helmet-mounted 360-degree 3D visual and audio system coupled with tactile input-output would enable pilots to "fly" simulated aircraft, instead of seeing just a narrow-field 2D display. It would allow the battle commander to move about a virtual battlefield and observe how well the virtual aircraft and other forces are performing. It would allow the system developer to enter the virtual environment and experiment with system parameters.

The most visible M&S technology needs involve improving the power and realism of Air Force M&S's by applying the constituent technologies discussed above. Particularly significant technologies in this regard are high performance computing and communications, software, human-computer interaction, and artificial intelligence.

However, it is also critical for the Air Force to address several less flashy but equally important M&S areas, which require a mix of technology advances and institutional emphasis. Primary among these are M&S interoperability; robust M&S; M&S verification, validation, and accreditation (VV&A); continuous instrumented exercise; and virtual system acquisition.

9.5 Modeling and Simulation Interoperability

There are many sources of M&S interoperability problems, involving incompatibilities in data definitions, control structures, timing, operating assumptions, and COTS componentry. These need institutional attention (e.g., via M&S adjuncts to Air Force Horizon interoperability initiatives) and new technical capabilities. The most significant technology areas are domain architectures for simulation and application domains; wrapper, mediator, and intelligent agent approaches for making M&S components interoperable; and M&S interoperability test and assurance technology.

9.6 Robust Modeling and Simulation

The more complex Air Force models and simulations tend to emphasize ad-hoc point solutions for demonstrations and individual exercises. Continuous two-sided M&S exercise requires significantly more robust M&S capabilities. The primary technology areas for M&S robustness are domain architectures and test and assurance technology. As much as anything, however, robustness needs institutional attention to the more rigorous applications of known software engineering techniques to M&S initiatives.

9.6.1 M&S Verification, Validation, and Accreditation

VV&A is critical to the credibility of M&S as a basis for making operational and acquisition decisions. An additional challenge is the VV&A of artificial intelligence techniques employed in M&S agents and semiautomated forces. Key technology opportunities include the development of more powerful and scaleable M&S validity assertion formalisms and checking techniques; and the use of intelligent agents for monitoring M&S validity conditions (although there is a bit of an agent-validity recursion issue here).

9.6.2 Continuing Instrumented Exercise and Virtual System Acquisition

Another approach for improving M&S interoperability, robustness, and VV&A involves continuing instrumented M&S exercise. Continuing exercise also strengthens M&S capabilities for training and decision support.

A particularly attractive application area for continuing instrumented M&S exercise involves virtual system acquisition. Providing capabilities for selectively mixing real and virtual system elements, and subjecting them to continuing exercise, can significantly improve the cycle time and effectiveness of traditional acquisition approaches. Critical success factors include domain architectures enabling modular mixing of real and virtual system elements; testbeds with instrumentation and analysis capabilities; revised acquisition process models and guidelines; and integration of technology M&S with life-cycle cost M&S.

10.0 Software Development

Introduction

The technologies addressed in New World Vistas are primarily focused on ensuring that on any given day in the future, say July 4, 2025, the Air Force is better prepared to prevail in combat than its adversaries.

However, suppose that on July 5, 2025, the Air Force detects that an enemy has discovered and is ready to exploit a critical weak link in Air Force combat capability. What technology can best enable the Air Force to repair the weak link and disseminate the fix to its full complement of forces?

An electronically disseminated combat system fix would best fit this need. The most powerful and flexible electronic fixes by far are changes to the combat system's software. Enabling the Air Force to design, develop, test, and deploy these fixes is the province of software development technology.

As with other information technologies, software development technology is increasingly paced by commercial investments. From an information warfare standpoint, this must be considered a technology leveler with respect to the Air Force and its adversaries. Not using commercial software development technology would be a competitive disadvantage for the Air Force, but so would a purely reactive use of commercial technology.

This chapter addresses Air Force investment choices, and their likelihood to provide competitive advantage to the Air Force through a pro-active approach to software development technology. It begins with a summary of the major leverage areas and elements of software development technology, along with likely future trends and Air Force implications. It then elaborates on the current Air Force software development status: a mix of complex challenges and significant initiatives. In this context, it then discusses primary opportunity areas for pro-active Air Force investments in software development technology. Finally, it discusses some key research and technology strategy elements which need rethinking for effective Air Force investment in a commercially intensive technology field such as software development.

10.1 Technology Definitions and Leverage Areas

Software development technology consists of methods, processes, tools, and assets enabling faster, cheaper, and better development of computer software. Methods, such as object-oriented design, configuration management, and Cleanroom methods, enable improvements in software design, verification, version control, and quality assurance. Processes, such as incremental, evolutionary, spiral, design-to-cost/schedule, and product line management processes, enable more efficient orchestration of the methods. Tools, such as design, code generation, test, product and process management tools, can completely or partially automate software development functions. Reusable assets, such as specifications, plans, components, and test cases, enable more of a software system to be composed of already-developed program modules.

Faster software development will be particularly important to the Air Force in an information warfare context; it will enable the Air Force and other US forces to "turn their information

systems within their adversaries' turn radius." Cheaper software development will enable the Air Force to contain its escalating software costs; for example, 30% of F-22 EMD-phase costs were software costs. Better software development will, in particular, enable USAF software to achieve greater-than-commercial-grade quality attributes where necessary (e.g., for performance, security, survivability, interoperability, and scalability).

By 2005, powerful user-programming capabilities will enable over 50% of the Air Force workforce to harness packages with millions of lines of code in a few hours to support a special mission need. Thus, for straightforward applications, the Air Force can get the software it wants right away. Getting all of these special applications to interoperate will require new layers of interoperability conventions. The resulting distribution of information processing power will combine with other trends to flatten Air Force organizational structures.

By 2005-2015, advances in software architecture technology will enable concurrent engineering of hardware devices and their software. This enables systems to have better allocation of their functions to hardware, software, and people. It also helps to get software off the critical path in the product cycle. By 2005-2015, semiautomated capabilities for determining the requirements and architecture of complex information-intensive systems will provide even more dramatic improvements.

By 2005, adaptive software systems will be emerging which automatically improve themselves based on observation of usage and data patterns. This will provide the Air Force with rapid, cheaper, more assured adaptation to changing situations (e.g., to adversaries' mobilization patterns). As with other adaptive systems used in information warfare situations, this requires defenses to avoid spoofing the adaptive system. (Techniques for automatic adaption are discussed in Chapter 7, Artificial Intelligence.)

Much of software development technology will be developed and supplied commercially, but some Air Force needs involve Air Force investments to achieve faster, cheaper, and better Air Force software. Examples are capabilities which scale to very large, complex, real-time systems; tools with limited commercial markets, such as weapon system software test simulator/stimulators; and reusable software asset generation for Air Force domains. The bottom-line result of this investment can be a continuous, closed-loop software upgrade capability, enabling the Air Force to adapt to dynamic changes in threats, environments, and technology at a faster pace than its adversaries.

10.2 Current Air Force Software Development Status

Future Air Force roles, missions, and environmental trends involve major challenges for software development technology. Some of these challenges will be well covered by the pace of commercial technology. Others will require particular Air Force attention and investment to ensure that the Air Force has a competitive edge.

The primary challenge areas for Air Force software development are affordability, combat performance, interoperability, information warfare, high assurance, and legacy software. Besides their individual challenges, these areas interact in complex and even more challenging ways. Each area is discussed below, along with its cross-area interactions.

10.2.1 Affordability

Decreasing DoD budgets, broadening threat spectra, and increasing commercial-off-the-shelf (COTS) software capabilities imply that the Air Force and DoD can rarely afford their previous luxuries of full-custom software. For the future, Air Force software development must treat COTS software as the primary driver of its information processing capabilities.

In Air Force-specific applications areas (combat aircraft, missiles, sensor processing, etc.), the Air Force can no longer afford its previous luxuries of stovepipe software systems with redundant capabilities. Instead, the Air Force needs to develop software product line architectures for its families of applications, and to reuse software assets across each product line.

10.2.2 Combat Performance

A totally COTS-driven approach to Air Force software would leave Air Force warriors with no competitive advantage with respect to adversaries, who will have access to the same COTS capabilities. Instead, the Air Force needs to identify areas in which better-than-COTS performance will provide significant combat advantages, and to invest proactively both in developing Air Force-unique software capabilities and in influencing the COTS marketplace toward compatibility with and support of Air Force software assets.

Particularly important software areas with respect to combat performance are embedded real-time software, high-performance sensor processing, intelligent software, and distributed information management enabling anywhere-anytime access to decision-critical information. To ensure the best “defense against tomorrow,” this software also needs to be rapidly and reliably modifiable, including warrior-tailorability wherever feasible.

10.2.3 Interoperability

In developing product line architectures, the Air Force needs to ensure that these do not become macro-stovepipes. Thus, Air Force software development and product line management needs to operate in the context of an Air Force enterprise architecture ensuring interoperability of combat systems. This enterprise architecture especially needs to evolve proactively with respect to the evolution of COTS software.

10.2.4 Information Warfare

The need to prevail in information warfare situations places particular stress on software development for the capabilities above. Not only does the Air Force need rapidly-modifiable adaptive software providing anywhere-anytime information to its own warriors, but also it needs to develop this software in ways which ensure its integrity and non-compromisability, and which selectively deny its availability to adversaries.

10.2.5 High Assurance

Implicit in the challenges above is the further challenge to provide high levels of assurance that Air Force software is secure, survivable, safe, reliable, interoperable, and scalable to high performance for large systems in crisis situations. These levels of assurance frequently exceed those considered adequate for commercial marketplace competitiveness.

10.2.6 Legacy Software

Further, the Air Force does not have the luxury of developing all of this software from a clean slate. The Air Force owns one of the world's largest inventories of antique legacy software running its current operations. Thus, attractive software development strategies such as incremental and evolutionary development are seriously constrained by Air Force needs to ensure continuing combat capability.

10.3 Current Air Force Initiatives

Several current Air Force and DoD initiatives are addressing portions of these challenges in effective ways. Recent acquisition initiatives are focused on replacing government bureaucracy and micromanagement by best commercial practices, enabling improvements in affordability and cycle time. Several Air Force and ARPA initiatives in domain specific software architectures (DSSA's) have created success stories with demonstrated returns on investment, particularly the AF/ESC PRISM initiative and the AF/ARPA STARS initiative for Cheyenne Mountain software.

The AF/SC Horizon initiative to develop and sustain a living Air Force software enterprise architecture has made significant progress, in concert with such DoD initiatives as the Technical Architecture Framework for Information Management (TAFIM). The DoD MIL-STD-498 effort, now culminating in IEEE standard 1498, will successfully harmonize, modernize, and replace two outdated standards for developing mission critical software (DoD-STD-2167A) and corporate information management software (DoD-STD-7935A). Air Force use of the Software Engineering Institute's Capability Maturity Model and related software capability models in source selection have significantly reduced risks of selecting immature software contractors.

These initiatives need to be strongly supported and extended to other Air Force domains; otherwise, Air Force software development will be increasingly cumbersome and constrained by its legacy software and culture. However, to be pre-eminent in the aerospace-cyberspace world of the future, the Air Force needs to develop a proactive investment and experimentation strategy to help shape and to capitalize on upcoming software trends. The main software development technology opportunity areas are discussed next.

10.4 Software Development Technology: Opportunities for the Future

Here are the primary opportunity areas in which Air Force investments at the margin in software development technology can provide significant Air Force competitive advantages.

- User Programming: the ultimate in rapid software development to meet user needs.
- Software Architectures: a key to user programming, as well as broader forms of rapid prototyping and rapid application development.
- Assurance Technology: analysis, development, and verification capabilities for the assurance of Air Force-critical software attributes, particularly security, safety, survivability, combat performance, interoperability, and crisis assurance.

- Software/System Concurrent Engineering: enabling integrated product teams to explore many software/system tradeoffs collaboratively, and to rapidly develop integrated software/system capabilities.
- Adaptive Software: software which can automatically improve itself based on observation of its usage and data patterns.

10.4.1 User Programming

For some applications domains, commercial capabilities are available which enable users to directly compose software applications in a few hours by specifying a number of options, parameters, and simple rules. Examples include spreadsheets, business fourth-generation languages, and tailorable packages for such applications as financial analysis, accounting, and inventory control.

Some similar capabilities are emerging in such Air Force areas as transportation planning, signal processing, and simulation. Air Force investments in further formalizing these and other Air Force-critical applications domains to create user programming capabilities will have significant payoffs in rapid user-responsive software development. However, the prospect of having many users developing special applications, and then discovering that many of them need to interoperate, also implies a need for counterpart investments in higher-level applications interoperability conventions.

10.4.2 Software Architectures

The primary key to user programming is the domain specific software architecture (DSSA), which provides the software structure for a family of software products in an application domain such as avionics, command centers, or signal processing. The DSSA determines a set of common application interface specifications, around which reusable software components can be composed. It also determines a set of common application concepts around which special purpose languages for user applications generation can be developed. It is based on a domain ontology, which organizes knowledge about the domain and its information processing aspects.

Besides application generators, which necessarily cover only a subset of applications in the domain, a software architecture provides support for rapid prototyping and composition of intermediate-scale applications, and for efficient development of large applications via reusable components.

Software architecture technology also provides complementary artifacts such as architecture definition languages and architectural styles. These capabilities are just emerging, but they strongly enhance the analysis of alternative software architectural options for such properties as security, safety, reliability, interoperability, and performance (The nature of a successful architecture is generally driven by such nonfunctional requirements; it is difficult to retrofit attributes such as security and fault tolerance into an application which did not consider them in its architecture). Strengthening these architectural techniques will help the Air Force to deal with its needs for critical nonfunctional requirements.

10.4.3 Assurance Technology

The ability to specify and analyze architectures for Air Force-critical nonfunctional requirements is one avenue of assuring systems with these attributes. Other complementary assurance technologies involve specification of desired attribute levels, including specification of the system's operating environment; and capabilities for developing, verifying, and evaluating high-assurance systems.

10.4.4 Software/System Concurrent Engineering

The Air Force is deriving considerable value from concurrent engineering of systems by integrated product teams (IPT's) involving multiple classes of system specialists and stakeholders (users, customers, developers, maintainers, interoperators). However, technology support for such concurrent engineering is still at a low level, particularly for integrated hardware/software engineering, large system development, and groupware capabilities for collaborative work.

Particular technologies that could enable more rapid and cost-effective Air Force software/system development are model-driven concurrent engineering (based on models of the domain, the system objectives, and the collaboration approach); hypercode (hyperlinking of software and system artifacts such as plans, specifications, code, and test procedures); and advanced processes and metrics for concurrent engineering (enabling both rapid execution and disciplined management of software/system development).

10.4.5 Adaptive Software

The artificial intelligence and agent technologies discussed in Chapters 7 and 5 respectively can also be focused on problems of adapting software to changing situations. For example, decision support systems could analyze patterns of decision information requests (exception reports, graphic displays, correlations) and restructure the software to better provide the most frequently requested functions. Such capabilities could both reduce software maintenance costs via automation, and effect software improvements which otherwise might not have been identified.

As with other adaptive systems used in information warfare situations, such software could be spoofed by patterns of requests for irrelevant or misleading information. This implies the additional need for defenses to avoid spoofing the adaptive system.

10.5 Rethinking Air Force Software Research and Technology Strategies

The software research and technology associated with a very large and increasing software marketplace requires the Air Force to rethink its own traditional software research and technology strategies. Focusing on Air Force software needs (e.g., for distributed processing) independent of commercial technology trends runs major risks of ineffective or commercially incompatible results. On the other hand, a purely reactive Air Force strategy would leave the Air Force with no competitive advantage with respect to other COTS consumers (and sometimes a disadvantage due to legacy software): a weak position from the standpoint of information warfare.

Thus, a new research and technology strategy is needed for Air Force software technology (and increasingly for other areas with growing commercial capabilities). For software, the major elements of this strategy include:

- Domain pre-eminence
- Living enterprise architecture
- Focus on Air Force-critical niches
- Pro-active commercial influence
- Close research-technology coupling
- Continuing closed-loop exercise and experimentation

10.5.1 Domain Pre-Eminence

Investments in software product line architectures, tools, components, and techniques for combat-critical domains such as sensor processing, data fusion, C3I, and combat platforms will give the Air Force a competitive edge with respect to adversaries possessed only with commercial capabilities.

10.5.2 Living Enterprise Architecture

Air Force combat interoperability requires that Air Force software product lines avoid becoming incompatible stovepipes. This requires a living enterprise architecture such as the Horizon initiative is addressing. Research and technology investments are necessary to keep the enterprise architecture in step with commercial technology—and further, that Air Force enterprise architecture needs are pro-actively inserted into commercial technology development and standards activities.

10.5.3 Focus on Air Force-Critical Niches

Besides combat-critical applications domains, the Air Force has needs which go considerably beyond those of mainstream commercial customers. The primary example area involves the high-assurance capabilities discussed above. Research and technology investments in these areas can also enable the Air Force and DoD to play an influential role in guiding commercial standards and technology.

10.5.4 ProActive Commercial Influence

Commercial products and standards will tend to focus on mainstream market needs, in which intermediate assurance levels will meet most customers' requirements. Often, however, commercial vendors will see a downstream advantage for incorporating more advanced Air Force and DoD needs into their next-generation products. A good past example is in software security, in which commercial operating systems, database management systems, and network software have been developed to DoD "Orange Book" security standards with the aid of DoD-developed security tools. To ensure this commercial relevance and influence, Air Force laboratories need to devote considerable effort to experimentation with and influence of advanced

commercial software technology and emerging standards. A significant current example is the Common Object Request Broker Architecture (CORBA).

10.5.5 Close Research-Technology Coupling

The rapid pace of commercial software technology means that Air Force software research and technology investments need to be well coupled with complementary software research across DoD, and also well coupled with non-DoD research and technology. Particularly important are experimental applications of new software technology to evolving Air Force software-intensive product lines such as those developed by PRISM at ESC and STARS at Space Command.

10.5.6 Continuing Closed-Loop Exercise and Experimentation

Besides development and deployment of new technologies, the Air Force needs a much stronger program of continuing exercise of its software-intensive systems, with feedback of weak and strong points to both the operational community and the technical community. This will be especially important for Air Force systems highly exposed to information warfare threats, in which feedback from continuing penetration teams can focus investments in software security, for example. This kind of closed-loop operation is another reason for the Air Force and DoD to evolve their contracting approaches toward long term customer-supplier relationships with strongly shared context and objectives.

10.5.7 Bottom Line

Changing Air Force software research and technology strategies will not be easy. But some of the initial Air Force success stories indicate that it is achievable. And with the revised strategies, when the tomorrows come, and the Air Force needs software quick-reaction capabilities, Air Force software development technology will be ready for them.

11.0 Software Infrastructures and Standards

Introduction

Computer software exploits the capability of computers and their attendant communications and enables all the applications to which computers are being put. The breadth of computer use in commerce permits the military to exploit the massive investment being made for commerce, manufacturing, education, and entertainment. This investment will continue to drive progress, and the benefits are available to all: commerce, the military, our friends, and our adversaries. While the Air Force cannot push progress in general software infrastructure capabilities beyond those boundaries, it must assure that its acquisition regulations permit it to remain current with the state-of-the-art.

This brief monograph will focus on issues where there are specific Air Force and DoD opportunities or issues. For a broad outline of status and plans for DoD software we refer to the proceedings of the July 17, 1995 workshop on "Defense Software Research, Development, and Demonstration: Capitalizing on Continued Growth in Private Sector Investment", by a panel of the National Academy of Engineering. That report draws in turn on the draft DoD Technology Plan for Software, developed by Barry Boehm et al., and on the analysis by the Defense Science Board Task Force on Military Software.

11.1 Definitions

Software infrastructure comprises the operating systems, communications, networking, input/output software, the file systems and databases, as well as many of the tools that are now expected to be delivered with a computer system. The scope of infrastructure software is increasing rapidly. As researchers or developers sense that some application function is repeated in several instances, it will be implemented as a sharable tool or service. If the tool or service is effective it will be marketed and in time displace functions previously bound to particular applications.

The type of tool will differ depending on the type of computer system. Mainframe computers will be equipped with transaction management systems for commercial processing, mini-computers for process control will have real-time data-acquisition and control software, engineering workstations will have support for shared access to design drawings and files, multimedia systems will have image-processing software, and personal computers will have word-processing, graphics software, and spreadsheets.

Services will be offered over the ubiquitous networks. Services include passive services, such as databases and other information resources which can be accessed by customers and their programs, as well as active services, such as software agents which will carry out tasks as searching and retrieval for specified types of data. Added-value services have an economic value, and are typically provided by commercial or publicly supported institutions. Mediators are processing services on the network which can access, transform, aggregate, and integrate data to provide information tuned to a specific type of customer, domain, or application type. Mediators provide added value, and are associated with an expert or an institution which takes responsibility for the content and quality of the result.

Software languages and their compilers are also part of the infrastructure, but have a reduced role today for the DoD compared with the days when DoD software was written in hundreds of languages and dialects with little concern for commonality with non-DoD practices. Today few principle languages remain, and they all have a specific niche:

1. FORTRAN for scientific computation and sensor data processing.
2. Ada for embedded and secure software modules.
3. LISP for artificial intelligence programs, especially their initial demonstrations.
4. C and C++ for systems and general purpose software.
5. Higher-order languages, which typically generate intermediate code in one of the languages listed above.

For the modern, distributed software systems the principal concern is the interoperation of software modules, often written in different languages and operating on distinct hardware platforms. While development and maintenance tactics will differ for different languages, the resulting software code can be indistinguishable.

11.2 Status of Software and Software Acquisition in the Air Force and DoD

Infrastructure software, as enumerated above, is making rapid progress in functionality, performance, and interoperability, driven by intense market pressures. It is not feasible for the DoD to advance the broad front of this arena beyond what becomes available Commercial-Off-The-Shelf (COTS), although there will be pockets where DoD investment can exert crucial leverage. There are invariably intersections where competitive pressures have a negative effect, for instance where vendors benefit from limiting interoperation with their competitors.

It will be difficult for the DoD to keep up institutionally with COTS infrastructure advances if traditional policies of global computer system acquisition, preceded by large scale evaluation and vetting, are followed. A careful, but slow process of determining equipment standards guarantees obsolescence. Two factors come into play here. The obvious one is the delay. By the time the DoD procurement system selects, approves and obtains new equipment, it is obsolete, relatively costly, and hard to maintain. It is rare that DoD computing equipment is less than 3 years old in terms of design, and will therefore necessarily be well behind the state-of-the-art. Having obsolete equipment means that modern software and existing infrastructure products are hard to use. Savings due to spending less on acquisition of hardware and software infrastructure induce major costs in applications and their maintenance, which would be avoided if systems were state-of-the-art.

The second factor is less obvious. Selling equipment to the DoD removes vendors from commercial realism and has caused their eventual demise. The DoD and Air Force acquisition focus must shift to defining standards for interoperation (as TCP/IP), and letting military units buy the best and most modern hardware and software that complies with those interface standards. Today the individual soldier often has a better personal laptop than the equipment provided by the DoD. There are numerous anecdotes from Desert Storm that illustrate individual initiatives where personal equipment provided solutions.

11.3 Projection for Investment

Not all applications needed by the Air Force will be satisfied by the commercial world. But even DoD specific software can benefit from using the best tools available. Using modern tools also helps in tapping the existing talent pool of software specialists, and reduces training costs. Given the rapid progress of the computer industry in providing infrastructure improvements and the widespread familiarity of recruits and contractors with computing, assessments of software investment are essential.

11.3.1 Progress Due to Commercial Interests

We focus on infrastructure software here: the operating systems and the basic network services that connect the computer systems.

Commercial interests are causing a rapid convergence to a smaller number of distinct software systems. While there were once dozens of mainframe operating systems controlling large computers, hundreds of intermediate systems, and a diversity of small systems, we can project that ongoing commercial investment will be limited to three to five families of basic operating systems, all available on multiple vendor platforms. We posit that there will be one operating system for mainframe computation, likely based on IBM practice, one for a family of intermediate systems, clustered around a standard operating system architecture based on UNIX or Windows NT, and perhaps two or three families of systems on personal computers. It is not unlikely that this number will reduce further, although there will continue to be vendor-specific versions. There will also be broad use of standards among those systems, for instance SQL to enable UNIX-systems to access mainframe databases, and CORBA to enable personal computers to access object-structured resources kept on UNIX-based servers.

These interoperability standards will continue to grow in breadth and capabilities, since, for each vendor, maximizing access to resources of others is crucial for winning market share. Similarly, features made available in one system will be rapidly replicated and superseded by features provided by a competitor. Specifically the arena of multi-media will continue to see rapid change, since the spread of visual information to millions of sites, driven by the entertainment industry, will motivate investment to lower cost and enhance capabilities. Even though the pace of change may slow down somewhat, there is little doubt that throughout the next 25 years the commercially available software infrastructure will provide the foundation for most DoD and Air Force systems. Even mission-critical applications will be supported in this manner, although the systems they acquire may be vetted for security and trustworthiness, causing DoD to lag behind the curve with respect to features and cost-performance ratios. This lag should be minimized, so that the differential of SW availability and its capability to exploit modern hardware does not lag excessively behind that of adversaries, who have fewer constraints.

Legacy systems will continue to exist beyond their economic lifetime because of conversion costs, but where the systems have to interoperate they will be wrapped to conform to existing standards.

11.3.2 Capabilities Appropriate for Government Support Funding

A concern for the national effectiveness of broad-based software systems is the developing of common and sharable systems, and the standards that enable interoperation, to avoid excessive dominance of a few vendors. This would lead to a concomitant isolation of computational communities and specializations, say operations researchers and knowledge engineers.

Vendor competitiveness motivates uni-directional compatibility, but of broader concern is bi-directional interoperation, so that users, and specifically DoD customers, retain the flexibility to change platforms and use systems that are composed with hardware and software obtained from multiple vendors, including competitors. Standards, and verification of conformance with standards, are tasks that should be supported by government, to assure fair and early dissemination of sharable technology. Such standards should focus in *interfaces*, not on equipment of software, so that best possible equipment can be obtained within interface constraints.

Infrastructure standards to consider include:

1. Operating system services as micro-kernels and POSIX
2. Base communications, as TCP/IP and successors
3. Database access, as SQL, RDA and successors
4. Object-oriented and multi-media access, as CORBA, MPEG, and successors
5. Knowledge transfer, as KIF and successors
6. Multi-representation interoperation, as OLE, KQML and successors

DoD can protect its interests by requiring its contractors to use standards rather than proprietary solutions whenever feasible, encouraging developers to use the most recent and emerging standards, and supporting researchers' participation and travel to the standards developing committees, so that the compromises of the vendors will be mitigated by forward-looking participants.

Within the US government NIST has the primary role for standards, including software standards, but has rarely been able to be proactive. There is hence a role for DoD research agencies to support standards activities that can accelerate the adoption and availability of sharable technology of interest to the military from commercial vendors. Examples of such involvement are the ARPA CSTO support of an operating systems micro-kernel, which can enable the transition of modern operating system technology into multiple vendors' military embedded systems, and the ARPA knowledge-sharing initiative, to assure interoperation and easy employment of artificial intelligence technologies into DoD information systems.

These subsystems, integrated into the commercially available infrastructure, will provide the sharable basis for future military information and control systems, and assure that software systems developed for the Air Force and other branches of the military can be flexibly and rapidly deployed. They will displace the bottom layer of the traditional stovepipes that characterize DoD system architectures. They will also provide the infrastructure for interoperation of general software development and the global information systems that are higher in the information system food chain.

11.3.3 Capabilities Requiring Military Investment

On top of the considerable commercial and national infrastructure will be some requirements that require specific military investment if US SW technology is to be maximally effective and superior to the technology of adversaries. Specific capabilities include:

11.3.3.1 Real Time

Highly robust real-time software modules are needed for data fusion and transmission. Real-time constraints require an ability to shed load or modify processing tactics in order to satisfy hard-time constraints on result delivery or feedback initiation. As improved fusion algorithms are being developed there is a need to incorporate them rapidly into the system infrastructure. While high-rate performance requirements are also common to commercial real-time systems, there is less of a need in commerce to assemble specific systems as rapidly as the demands of newly emerging threats. In the commercial world real-time systems, such as those used in refineries, factories, or power plants, the systems are carefully designed and evolve as the facilities are upgraded. Filters exist to assure that no unexpected signals enter the real-time processing stream. These conditions cannot be counted on in military situations.

11.3.3.2 Integration Of Simulation Software Into Military Information Systems

The support of decision making requires simulation, but the commercial tools lag far behind those used for advanced military simulation for training and modeling. The current applicable commercial interoperation standards focus on components such as spreadsheets and graphics, and permit at most dynamic incorporation of spreadsheet results into textual documents. Standards such as CORBA permit access to general object-oriented data, but little flexibility is provided, so that simulation objects cannot be easily aggregated, or partitioned.

It is doubtful that commercial progress relevant to interoperable simulation will support military requirements, if the planned insertion of simulation into all phases of military systems acquisition, testing, deployment, training, planning, and plan execution proceeds at the pace envisaged. At the same time, conceptual and structural similarity with developing standards in the COTS world should be maintained, so that plug-and-play capabilities are maintained. The Air Force must not put itself in an isolated position.

11.4 Paying Attention

Organizationally, it will remain critical for the Air Force Labs, such as Rome Lab, to maintain close linkages with ongoing technology and continuously monitor the gap between the available and emerging commercial technology and military needs, identifying the high leverage items that warrant the attention and funding of the Air Force.

12.0 Computer Hardware and Architectures

12.1 Circuits and Devices at the "Bottom"

The late Nobel Prize physicist Richard Feynman once wrote a provocative visionary paper that discussed the "room at the bottom." He predicted, indeed championed, a revolution in our understanding and manipulating of micro things.

For more than two decades, the computer world has been finding and using a great deal of room at the bottom. The lines that circuit makers engrave on their silicon (and other strata) have gone from many microns wide to quarter-micron wide. Consequently, the number of transistors on a chip has risen from the tens of thousands to the tens of millions. Since small also means fast (because of speed-of-signal-propagation considerations), these circuits have become many factors of ten faster. The story is familiar to all: next year's Pentium is 1.6 times the speed of last year's Pentium, barring the fact that memory speed doesn't keep up. But the net is still impressive with a conservative doubling every 3 years.

Not so familiar, except to consumers who may have the insight to wonder why they are getting their new PC disk drives so cheaply, is the revolution in storage density on rotating magnetic and optical media. Here the gains have been tracking the gains in the chip arena. Both have seen a doubling of performance per dollar every eighteen months or sooner. The machine that is being used to write this is three years old. It came standard with an 80 MB hard drive. The 1995 version comes with 500 MB.

The changes are revolutionary because they are exponential. With exponentials you can forget about the past. The main impact of exponential change is new structures, not the evolutionary change we've seen with the relatively constant priced PC of 1981.

Exponential change is revolutionary because it moves things from minute to immense in a short time. Our slow-moving minds, organizations, and infrastructure, accustomed to at most linear change, are overwhelmed.

Does the bottom have a hard floor somewhere? For years, people have been predicting a real floor for silicon circuit technology, and none of those predictions have come true. The latest prediction is that at about one tenth micron, circuit characteristics will become unstable and error-prone. The year may be 2003-2006. New materials will have to be found or new computer system architectures will need to be used to preserve the historical growth curve. These will be discussed later.

Magnetic disk technology is an exquisitely refined technology, driven to that state by intense competition in the industry. Magnetic bit density advance may stop its exponential growth within the next decade. Optical disk technology is not yet the replacement, in cost or performance, primarily because it has not yet received the major development that intense competition and high volume bring. Both of these technologies for secondary storage share a common architectural future: parallelism. Parallelism fits awkwardly into the traditional modes of software writing. But parallelism fits data base filing and searching very well, providing another avenue for continuation of the growth curve. We discuss this later.

12.1.1 New Materials

Beyond electronic switches, what phenomena in nature can provide reliable computing (i.e., behave in a “machine-like way”)?

A molecule of DNA can be thought of as “computing” the proteins whose manufacture it guides. Experiments have been recently reported in which a classical mathematical problem was encoded as DNA’s base sequences. A modest quantity of DNA (containing almost a million million DNA molecules) performed what is in effect an analog computation. The solution was retrieved by standard genetic engineering techniques. While primitive at this stage, this kind of molecular-level computation probably will open the door to a future of immense parallel computations. The search is on for other molecules that are effective “computers.” (See discussion in High Assurance Systems, Chapter 13.)

Physicists are exploring as another form of “analog computation,” certain phenomena that are implicit in the equations of quantum mechanics. The specific class of computations under study for the so-called quantum computers is factoring. Factoring might be considered to be a rather narrow and perhaps unimportant class of computations, until one recalls that modern strong encryption methods rely on the difficulty of factoring!

Molecular computing and quantum computing are long shots—long in odds and (will be) long in coming. However, the parallelism and speed they could bring would make even today’s revolutionary change seem evolutionary.

Some effort is being invested in optical computing—information processing on “bits of light” using “light switches”. There is also work on hybrid electro-optical computing, where the physical bit representation is moved from the circuit domain to the light domain and back, to capture the advantages of the technology of switching circuits. The gains in speed that are possible using light as the medium are potentially very large, but progress here is slow. We may not see optical or electro-optical computing during our forecast period.

12.2 Computers-on-chips and Their Architectures

There are two key concepts underlying predicting the future of computer systems:

1. The economics of volume production
2. Parallelism

12.2.1 Economics

Why is Intel like the New York Times? Is the analogy farfetched? No. Each produces its principle product by a kind of printing process. Each achieves low cost of the product by large volume production. The revolution in computers has been essentially a printing revolution (shades of the revolutionary technology of Gutenberg).

The costs associated with producing microcomputers are extremely large. These costs form a big barrier to entry into this business. Over the next decade, we will see the number of microcomputer makers shrink to a literal handful; and most of the microcomputers made will be made by Intel.

The cost of chips will range from several dollars (for smaller or older designs) to several hundreds of dollars (for the latest, fastest). Since these costs are small compared with the value of the end product, *many* will be used *together* in various configurations to satisfy the different “sizing” needs of computer architects. Think of them as the prefab modules of a semi-custom house.

Large numbers of the latest and fastest will be organized together on very-high-speed, in-the-box networks to form a supercomputer. Much smaller numbers will provide the computing power for an engineering or graphics-arts workstation (same network-in-the-box idea, however). One or a few microcomputers will inhabit the various types of personal computers we will have. Scaled computers will have common software and I/O compatibility.

12.2.2 Parallelism

Techniques of parallel processing inhabit all levels of computer design. Only the end-user is shielded (more or less) from their complexities. Parallelism is now one of the most important themes in university computer science departments and research labs.

At the chip level, microcomputer architects decompose and rearrange (with remarkable ingenuity) the basic instruction and communication processes to obtain as much concurrency of processing as possible.

As discussed earlier, the systems of the future will be networks of microcomputers built mostly out of merchant-standard components. The microcomputers will often be clustered, with very high speed communication within a cluster, and lower inter-cluster rates of communication.

The concept is very simple and the technology is likely to evolve to be relatively simple to do, once microcomputer chips are in hand. This implies that it will be relatively easy for potential adversaries to put together virtually any size computing engines that they need for command and control or for weapons systems control. Beyond military, the same is true of course in commerce.

What network architectures will be used in this style of architecture? ATM for LANs shows great promise, and experiments are being done using ATM for the internal “LANs” as well. But it will be many years (but not many decades) before the cost/benefit ratio for ATM is as low as today’s conventional networking technologies.

12.3 Pictures and Sound

The foregoing discussion was built around the concept of “computer” as we have known it for 50 years. Even the PC has been a kind of “little mainframe”. But the concept is changing. The computer of the next few decades will process bits at literally billions per second, but will not be doing much “computing” (i.e., numeric or symbolic). The bits will be digital video bits, digital audio bits, and the bits from sensors (see Personal Computing, Chapter 3 and Human-Computer Interaction, Chapter 4).

Handling pictures (with speeds and qualities that people want to buy) will be the “killer app” of the next few decades of “computing.” In the future Communication Sciences Departments

will replace today's Computer Science Departments. MIT has had one of these for a long time—the visionary Media Laboratory.

Our communications have been dominated by *telephones* and *television*. In the future the device will be the *telecomputer*. This is a fairly strong prediction. It says that the device to which this era of so-called “convergence” will converge is the computer-become-video rather than the cable TV set-top box (i.e., Microsoft, SGI, and Oracle rather than Time Warner, TCI, etc.).

For architects, this view of the future implies new architectures. The chips will be largely devoted to video I/O and sensor I/O. Most of the bits/sec flowing in or out will be representing pictures, not text.

A company like Intel is predicting that the majority of its revenues in the year 2000 will not be microcomputer revenue but rather revenue from a variety of chips to service video (and other data) streams! The fiber lines to homes and offices will enable this part of the information revolution.

The new “computer/communications” architectures are only now being worked out. So the design is not yet clear—in fact there will be an evolutionary competition for best surviving design. What role will parallelism play in the new designs? What will be the special features of the telecomputer architecture that will make it easy for the software developers to innovate and keep applications flowing?

12.4 Some Raw Performance Guesses

These guesses are “raw” in two senses:

- a. In IT, it is difficult to make quantitative predictions too far out.
- b. As important, the concepts themselves are “raw” as we move away from familiar territory to the new systems. For example, the traditional measure, Megaflops (concerned with the speed of calculation), makes less and less sense tomorrow as the percentage of bit-processing devoted to calculation decreases drastically.

That said:

1. A billion operations per second (BOPS) in desktop workstations in 1996-97 (expensively) and 2000 cheaply. Unfortunately, using standard “rules of thumb” one gigabyte of memory is needed to keep the machine balanced. Unless we see a dramatic change in memory pricing, such a machine seems impractical.
2. Four BOPS in high performance workstations by 2000, PCs by 2002-3.
3. Supercomputers using a “network of workstations” architecture achieve a thousand billion OPS (TeraOPS) by 2000 or sooner. Such machines have a commercial and economic realization by 2010. What will hold up such designs is the relatively small demand for such large machines. Microsoft is dedicated to creating an “upsizing” environment based on clusters of commodity PCs and ubiquitous networks such as ATM.

4. Having defined the end-points of the spectrum, almost any capability in between will easily be possible because of the plug-and-play nature of network of workstations (NOW) architectures (discussed earlier). The marketplace will abound with options for customers.
5. In 10-15 years digital video on "telecomputers" will be as common as television is today.

13.0 High Assurance Systems

13.1 There Are (as yet) No Rules of Engagement In Cyberspace

A great deal of attention has been given to the notion of information warfare, with considerable use of terminology and metaphors from more traditional forms of conflict. But there are also a number of ways in which conflict in cyberspace is importantly different. As one example, we will at times not know whether we are under attack, because information attacks can be considerably more subtle than physical attacks. Information attacks can occur from arbitrary physical distances and involve no movement of personnel or materiel.

Even more important, there is not yet any established notion of *rules of engagement*; simply put, there is no agreed on notion of what constitutes a hostile act in cyberspace. Some possibilities are of course relatively obvious, e.g., breaking into password-protected systems. In some cases state law has begun grappling with this via legislation concerning "computer trespass." But even here there is some ambiguity in distinguishing a break-in attempt from an innocent mistake. There is also the difficulty of issues of jurisdiction: given the arbitrary distance involved, acts that may be illegal in the US can be undertaken from other countries. While this phenomenon is not new, this does underscore the need for agreement that is as widespread as possible, ideally global in scale.

As with rules of engagement in other forms of warfare, the crucial properties of rules for cyberspace would be: (a) they are widely agreed on and (b) they are technologically feasible. As such the task of creating such rules is a mixture of policy and technology requiring participation from both communities.

Near-term steps:

The Air Force should assemble a high-level policy and technology task force to formulate draft rules of engagement.

13.2 Cryptographic Coding Will Be Unbreakable; Systems May Be Breakable

Cryptographic techniques currently exist in COTS that can provide very high levels of security to individual systems. *Public-key systems* in particular offer an encoding technique that has withstood considerable testing and appears to be unbreakable. By "unbreakable" we mean keys can easily be provided which ensure any attempt to decrypt an encoded transmission would require arbitrarily great effort. One can, for example, use keys which ensure breaking the code would require tens of thousands of years of computation on machines tens of thousands of times faster than anything available. If that is insufficient, a few more digits can be added to the key to make the effort required hundreds of thousands of years, and so forth.

Over the next 5 years integration of this technology into Air Force software at all levels (networks, operating systems, and applications) will provide security against any attempt to extract information by decrypting intercepted signals.

Important attention should now be paid to Air Force policy regarding cryptology. In particular: 1. as a COTS technology, the major stumbling block to deployment and integration

will be policy rather than technical. Appropriate authorities must permit the use of these systems. 2. we recommend strongly the Air Force employ a key escrow (or similar) system, in order to ensure that internal use of cryptographic techniques cannot provide an impenetrable wall of privacy to unauthorized action by Air Force personnel.¹

Formulating appropriate policy and effecting deployment of cryptographic technology are important near-term steps because existing information intensive systems are currently blatantly vulnerable. Recent studies have shown the domestic electric power grid, major financial systems, and the telecommunications infrastructure to have between modest and virtually non-existent protection against information-based attacks. Break-ins have occurred via techniques as primitive as knowing the phone number of a modem at a site. Military systems have routinely displayed only slightly better protection.

There is no reason for this to be so; basic techniques exist to defeat many of these attacks, systems simply need to be designed to take advantage of what we already know how to do. Continued vigilance and development of counter-measures will of course also be necessary, as new modes of information-based attack will no doubt be created.

Information systems will continue to be vulnerable. Code-breaking is only one of the routes to the cleartext of an encoded transmission. Unbreakable codes offer no new protection against traditional techniques such as traffic analysis or compromising personnel. In fact, we anticipate that with the widespread use of unbreakable codes, attack effort will shift markedly, away from decryption and toward other approaches. This likely reallocation of effort should be considered when developing security policies.

13.3 In the Long Term, Low Probability Events May Impact Encryption

Current techniques have been well tested, but their ultimate status as unbreakable is contingent on certain mathematical hypotheses that have themselves been widely examined, but are not yet proven. One such hypothesis is that factoring large numbers is an inherently time-consuming process; very low probability discoveries may yet prove this untrue. A more general hypothesis in computer science goes by the name of $P \neq NP$; if this turns out to be false, a new basis for unbreakable codes will have to be found. Note these are extremely unlikely events, and even if they were to occur, their initial significance would be strictly theoretical.

Two other lines of development should also be monitored. New models of computation have recently been proposed and explored in a very limited way. One approach—molecular computation—involves the use of sequences of DNA as a way of searching very quickly through a vast number of possibilities. It relies on the ability to create on the order of 10^{17} copies of a molecule in a solution, and the ability of one DNA molecule to bind to another that has a specific sequence of sub-units. It thus provides in effect 10^{17} extraordinarily simple computers that search

1. Key escrow systems involve keeping a copy of cryptographic keys with one or more trusted authorities. Procedures are established by which the escrow agent can divulge a key to appropriate personnel only after the required degree of need and authority have been established. Keys can also be split into two or more fragments, no one of which is sufficient to break the code. The fragments can then be parcelled out to a like number of authorities; in that case no one authority acting alone can break the code.

in parallel for an answer. A second approach—quantum computing—is even more speculative, involving the use of quantum mechanical effects to do computation. While both of these are low probability long-term developments, it will be useful for Air Force personnel to monitor their potential developments over the long term.

Near-term steps:

The Air Force should develop and deploy a key escrow system that is a routine and transparent part of any use of unbreakable advanced codes.

The Air Force should integrate off-the-shelf cryptology into software at all levels (networks, operating systems, and applications).

13.4 Biometric Identification Will Be an Embedded Technology

Biometric identification—identifications via physiological traits such as face recognition, finger-, retina-, or voice-print, are all currently effective laboratory demonstrations. We believe in the next five years they will routinely be embedded in devices and will provide useful levels of security by fusing results from multiple sources: voice, retina, thumb, and face recognition in combination can provide useful levels of performance, including acceptable levels of false negatives, and are difficult to counterfeit simultaneously.

We believe within ten to fifteen years biometric identification will be unobtrusive, continuous and ubiquitous. We predict a world in which “smart locks” can be embedded everywhere control is required. Those locks will be made unobtrusive (i.e., check identity without disturbing the user) and continuous (constantly verify user identity), thereby providing a high level of security at all times.

13.5 Survivability and Assurance Is Significantly More Difficult in Large-Scale Distributed Systems

Difficult problems arise in providing high assurance and survivability for larger-scale distributed systems. A significant body of foundational work exists in fault tolerance for modest scales of distribution, but significant work needs to be done to provide assurance and survivability when potentially every computer in the Air Force will be interconnected. Of particular note will be development of techniques for *graceful degradation*—the ability of a system to provide dynamically selected partial functionality in the face of unanticipated failures, rather than the all-or-nothing functionality provided by today's approaches to fault tolerance.

One interesting example of survivability is the Internet itself: over the past twenty years it has proven to be a remarkably robust means of communication despite a complete lack of centralized control. Interesting design lessons emerge from considering how and why this is so, lessons that may be of use in future systems. The net achieves its robustness by massive dispersal of functionality—tens of thousands of computers handle the message routing function. As a consequence, it may be difficult to cripple the net: it is too big and too widely dispersed. But that same broad distribution for routing also provides greater access to users of good and bad intent and any broad homogeneity may also be a vulnerability.

One possible reason the Internet itself, as opposed to its attached hosts, has remained relatively unthreatened and intact is that it is often the only means for a network terrorist to access a target host or server. To attack the Internet's routers or gateways would be to remove the conduit through which damage is perpetrated or information sought is recovered. Should one simply want to deny others use of the Internet, then it might not appear so invulnerable. Attacks on routing tables, as an example, could perhaps deny selective segments of a network, leaving other segments in use. This type of vulnerability must be dealt with for any network that becomes part of the AF infrastructure.

The research challenge for the future is to develop design techniques that permit us to provide more complex behavior in a similarly dispersed (and hence survivable) manner.

Near-term steps:

Support work in and through AF Labs to develop the ability to design systems with the ability to degrade gracefully in the face of unanticipated failures.

14.0 Collaborative Computation

Introduction and Definition

“Groupware” is the label for the notion that computers and information can be used to support business teams, rather than mere individuals (the dominant personal computing paradigm in the last decade). Groupware thus encompasses both the creation of entirely new systems, as well as the pressing into group service of systems designed for individual use.

PC-based groupware products are proliferating today, yet the term remains more a concept than a product category in its own right. In the long run “groupware” as a discrete category will disappear entirely, as group-oriented features are incorporated into virtually every product offered in the PC sector.

Groupware as a product category is growing solidly but modestly. The single largest groupware product is Lotus Notes, which is enjoying considerable popularity among corporations, but still has an installed base which is minuscule compared to the total installed base of networked PCs in corporations. The features that make it popular will find their way into more generalized environments such as operating systems, advanced data repositories and even Internet browsers and servers.

Groupware will have its largest impact on organizational structures. Even the most casual glance at business history makes it clear that each time a new infrastructure becomes available (e.g., railroad, telegraph, telephone) the entities which are ultimately most successful are also the first to reshape their structures in order to gain maximum advantage of the new information conduits. The new networks emerging today are “geodesic” (a term first noted by Peter Huber in the mid-80s in the context of telephone deregulation), that is, global, non-hierarchical and without any central node. It is a safe bet that our organizations will follow suit.

Ultimately, the term “groupware” will disappear sometime within the forecast period, for the simple reason that everything in computers and networking is becoming group-oriented. A decade ago, the term was a novelty in a world where “personal computing” was synonymous with stand-alone systems. In 1995, personal computers are anything but standalone, running with operating systems designed for network and Internet connections. In a decade or so, group functionality will be the norm and stand-alone anything will be the exception—perhaps it will be called “soloware” to distinguish it from the new majority group-oriented applications.

This group-oriented future will be relatively slow to arrive, but the Air Force can gain much by aggressively pushing group-oriented systems. In a sense the Air Force is already on the bleeding edge of groupware by virtue of distributed simulation operations on the Defense Simulation Internet.

14.1 The Laser: Groupware’s Key Driving Force, and Its Consequence for Communications

As mentioned earlier in the Personal Computer section, the advent of cheap lasers is the defining technological force in this decade. In turn, it is triggering a profound shift in the nature of communications as a medium, and it is this shift that groupware is riding into everyday reality. Decades worth of cherished telecommunications dreams may finally be realized, but

they are likely to come to pass in utterly unexpected ways. Communication isn't just coming into its own; it is becoming an entirely new medium. Three profound shifts can already be seen:

14.1.1 Communications: No Longer People Talking to People

For most of this century communications has been synonymous with people talking to people. However, the future growth in communications lies not in people talking to people, but in machines talking to machines on behalf of their human owners. Hints of this trend already abound. For example, voice conversations account for less than half the traffic on AT&T long lines between the United States and Japan the balance is generated by fax machines sending documents back and forth on behalf of their owners. The sales of fax machines are growing explosively.

The explosive growth of Internet use is another example of this shift towards a world of machines talking to machines on people's behalf. Email is the most widely used groupware application today. And this growth of communicating machines has already had unexpected impacts—the collapse of the North American Numbering Plan into a mess of new area codes was occasioned by the unexpected need to give machines their own phone numbers.

But so far, the devices communicating on our behalf have been helplessly stupid. Who trusts their fax machine enough to ignore the status printout? Add processing power, however, and things get interesting. This on-board intelligence will allow our communications devices to become more reliable and more autonomous, eventually becoming powerful “infobots” managing a growing list of more sophisticated tasks on our behalf, from managing calls and mail, to coordinating our calendars to purchasing information and conducting simple transactions. Ever more autonomous systems will of course create special challenges on the battlefield of the future. (See Intelligent Software Agents, Chapter 5.)

Simple examples of this emerging world already exist today. “Program trading” on Wall Street has been commonplace for years. The Internet universe is inhabited by a menagerie of simple software agents performing routine tasks like mail handling and information gathering for their human owners, and the first telephone agent, a system called “Wildfire”, is available for purchase. We will see a dramatic increase in machine-mediated communications in the groupware arena over the next decade. Look in particular for intelligent group schedulers, communications coordinators, group-oriented project and mission planning aids, and of course, ever more autonomous network-based agents toing and froing in the service of team-oriented masters.

Of course, new worries will replace old as communications become more machine mediated. We will wish for the simplicity of jammed paper as we worry whether errant software agents have misdelivered our electronic documents or are off making mischief somewhere on the global network. Larger-scale snafus will also be a certainty: recall that unsupervised software trading programs were what sent the stock market into a nose-dive in 1987. The key challenge for knowledge workers in this new regime will be one of intervention keeping up with the inter-machine conversation sufficiently to catch and correct glitches. Defining the limits of infobot autonomy and teaching our new infobot companions when to call for help will be a hot issue early in the next century. Groupware in the decades ahead will involve more than human-to-human conversations—“participants” will include machine intelligences as well.

14.1.2 Communications: No Longer a Conduit

A second shift goes even more directly to the nature of communications as a medium. For most of this century, communications have been a conduit, a pipe between distant physical locations. A decade or two from now, our communications medium will seem less a pipe than a location in its own right. It will become a place, a destination where we will conduct more of our business and personal interactions.

This future has been richly anticipated by today's cyberpunk authors like William Gibson, who describe a future "cyberspace" where we will spend ever more of our lives, and reality is quickly catching up. The popularity of MUDs, multiple-user dimensions, on the Internet hints at the shape that communications-as-destination will take. Despite the limitations of the text-only environment imposed by today's networks, the current crop of MUDs has matured into social virtual realities offering a richness of interaction unmatched by anything short of face-to-face encounters. The first video-based MUDs have been developed in several laboratories, and seem to be clear harbingers of the form that net-based systems will take over the next two decades.

Cyberspace today is largely monochrome and text-based, but it will evolve into something that is steadily more vivid, thanks to growing communications bandwidth. The first multimedia MUDs already exist. MUDs employing two-way video are unremarkable fixtures in a handful of research labs, and consumers will shortly be able to subscribe to a next generation of MUDs that substitute cartoon-like multimedia worlds for ASCII environments used today. In the long run, MUDs will mutate into more generalized environments where participants will be able to do the sorts of things people do in the real world today: socialize, hang-out and conduct business. Cyberspace will thus become a powerful adjunct to today's real-world business environment.

14.1.3 Communications: From Scarcity to Abundance, and Ultimately, Ubiquity

Finally, the communications medium is experiencing a third shift, from scarcity to abundance, and ultimately, ubiquity. This shift has already profoundly affected the regulatory environment which, since 1970 has been shaped above all by the consequences of communications abundance. Deregulation got a push from politics, but it was also inevitable, for the old regulatory order was preoccupied with the fair use and allocation of a scarce resource communications bandwidth or monopoly over wire connections to the home or business.

It is the final dimension of this shift, from abundance that ubiquity that will deliver the greatest surprises for planners. No matter how quickly communications reality advances, user expectations will advance more rapidly yet. Policymakers and soldiers alike will expect communications to be instantaneous.

14.2 Groupware: How Soon, How Fast

Groupware is growing slowly in terms of general business acceptance, but the Air Force has an important opportunity to take advantage of this technology. The combination of the domain-specificity of the Air Force mission, combined with its unique management structure positions the Air Force to take early advantage of emerging groupware system options, at a rate more rapid than is possible for business at large. With this in mind, this section offers a rough

forecast of which areas are likely to advance most rapidly to a point where they can be effectively exploited.

14.2.1 Videoconferencing

Videoconferencing systems are slowly emerging from a long gestation period that dates back at least to AT&T's PicturePhone service of the early 1960s. Videoconferencing rooms have given way to videoconferencing carts and most recently, to "desktop videoconferencing" built into workstations. The price of systems continues to plummet, both for terminal equipment, as well as communications links. In addition to the evolution of traditional commercial products, such as Intel's ProShare, two-way video is burgeoning over the Internet with the diffusion of Cu-C-me and M-bone. At the moment these Internet solutions deliver images little better than that of TV in 1952, but they are evolving rapidly, and are likely to deliver satisfactory video for ordinary users by or before 2000.

Diffusion of videoconferencing systems has been modest to date. Even low-cost desktop video systems are moving slowly. However, the situation is akin to that of fax machines in the last decade; the industry seems poised for a take-off sometime within the next decade. All that is missing is some stability in standards and more coordination among manufacturers. This will happen, though the moment of take-off of course cannot be predicted with certainty.

Plain-vanilla videoconferencing can be an effective business tool, but it is likely that the "phone call" model of setting up conferences will erode in favor of more socially rich and interactive modes of communications over the latter half of the forecast period. This represents a convergence of videoconferencing with novel groupware notions. Look for "video hallways"—permanently established links between informal spaces in physical sites separated by distance. Also, we will see the emergence of video MUDS that offer greater opportunities for informal interaction than formal video conferences.

This is an area where the Air Force can take early advantage of emerging systems. Businesses on the outside will have to wait for costs to drop further and for systems to be acquired by their client communities. The Air Force however can accelerate deployment of emergent systems into mission-critical areas as needed.

One footnote: beware of the "travel substitution" mirage. The simple fact is that rich electronic communications lead to the desire for more face-to-face meetings, and more face-to-face meetings inevitably lead to more electronic interaction between meetings. Instead of travel substitution, the Air Force should pursue "travel shifting"—the use of electronic communications in order to be able to cluster and rearrange physical meetings for greater convenience and effectiveness. Pursuit of mere travel substitution will amount to the sacrifice of effectiveness in favor of simple efficiency.

14.2.2 Shared Databases/Group Memory Systems

Lotus Notes is the current leading product in this area, but it is likely that we will see the emergence of competitive systems built around more open Internet-based models early in the Forecast period. Notes is likely to have a long life, and be upgraded by its new owner, IBM, and it will find a solid home in a small and slowly growing number of large corporations. But the more open systems are much more likely than Notes to capture large audiences outside of these

corporations, and ultimately will eclipse Notes in much the way Internet-based mail and the World-Wide Web eclipsed proprietary on-line systems.

One can forecast this broad trend with a high degree of certainty, but the details defy prediction. That said, the trajectory of surprise is likely to match that of the Internet over the last three years, and the requisite innovations are likely to come from the same quarter—university researchers establishing start-ups to exploit new markets.

14.3 MUDs

Multiple User Dimensions—MUDs—represent the fastest growing sector of the Internet today. MUDs are in effect social virtual realities that for the moment are largely text-based, employing a simple Pascal-like language to manage the details of the social interaction occurring. The majority of MUDs are recreationally oriented, but the number of research MUDs is increasing, and the first business MUDs are likely to appear in the near future as well. As already mentioned, multimedia and video MUDs exist in research laboratories today, and will begin to find their way out into the real world soon after 2000.

The main event on MUDs is human-to-human interaction, for the moment at least in the virtual MUD space, the humans “meet” in social, business, or gaming context. However MUDs are natural environments for the deployment of agents and near-agents. As agenting technology advances, human-agent interaction could become even more important than pure human interaction.

14.3.1 New Operating Systems

Operating systems have become steadily more network-oriented and group-oriented. For example, two major operation systems of today, Windows and the Mac OS, both include network access and file sharing as standard elements. As operating systems advance, look for the incorporation of email and other group functions as essential features.

One wild card is the disappearance of operating systems, replaced by network environments. In effect, the environment becomes so group-oriented that the very notion of an operating system begins to evaporate.

14.4 Implications

The growing diffusion of groupware presents a large number of implications for the USAF. A few of the most important include:

- Continued change in organizational structures. Our organizations have long been shaped by the communications infrastructures they use, and the 1990s will be no exception. As new communications media proliferate, look for new organizational forms to emerge as well. Hierarchies have already yielded to team-based forms; new web-like organizational structures appear to be the logical next stage. By the end of the 1990s, we may discover that laser-driven access has blurred organizational structures and boundaries to the point where the corporate model as we know it becomes all but obsolete.

- Electronic marketplaces and electronic commerce. Electronic commerce is merely a specialized form of groupware. The electronic trading systems of the 1980s amount to primitive predecessors to new venues for electronic commerce. Already, less than 25% of the US money supply is represented by paper currency. Now business *people* are joining their transaction flows in this new environment. Look for future hybrids of trading, conferencing, and e-mail systems to mature into the cyberspace version of the medieval marketplace, a transactional venue for the conduct of everyday business. The long-term consequence may be the emergence of “electronic commerce” the creation of forms of transaction as different from today’s modes of interaction as today’s business interactions are from those of half a century ago.
- Access tools will deliver on early 1980’s visions. For example, powerful workstations could finally make telecommuting a reality for more than a determined few enthusiasts willing to put up with the complexity and aggravation of PCs. Elsewhere, the diffusion of connectivity-rich information appliances is likely to support a whole new class of nomadic executives relying on a global network to work not at home, but everywhere.
- New generations of office tools will appear. Tasks, once performed by discrete devices like fax machines and copiers, will be combined and performed by a new generation of office tools. For example, companies are already beginning to offer inexpensive desktop “I/O utilities” that combine laser printer, scanner, fax, and low-volume copier in one printer-sized box. Workers will still go to the copy room to make large numbers of copies on the now-digital copiers, but these new boxes will save shoe leather when making one or two duplicates in the local office area.

14.5 Thinking the Unthinkable: The Virtual Pentagon

Imagine that groupware technologies become so effective that the Pentagon is replaced by a “virtual Pentagon” a destination in cyberspace accessible by Air Force personnel from wherever they happen to be on the planet. It would be premature to seriously propose such an outcome based on currently available technology, but it is likely that such an eventuality will seem a non-controversial option well before the end of the forecast period. In the meantime, we recommend using the notion as an organizational artificial horizon against which to measure the potential benefits of groupware advances. One extreme-case scenario follows.

January 2015

The last USAF officer has departed the Pentagon—and Washington. By now everyone in the Air Force is working much closer to USAF’s actual operations, dispersed at bases around the US and the globe. The Pentagon’s population has shrunk to that accounted for by the office of Secretary of Defense, and its staff, plus a small population of staff whose role it is to support visitors to the Pentagon. Military officers still visit the Pentagon, but they now “hotel” for brief periods when on-site for various meetings. Overall, the daily population of the Pentagon present on military business is one-tenth of the population in 1995.

The Pentagon's warren of offices has been radically revamped. Gone are the acres of office suites, replaced by a large number of flexible meeting spaces, and visitor offices. With slight variations for each service, everyone follows a similar pattern: an arriving officer stops first at a "concierge" desk, is handed a secure wireless phone and is assigned a visiting office. The art of the concierge is to cluster visitors together according to the reasons for their visit. For example, a team of 15 individuals has converged to meet around the development of a new weapons system. All are assigned offices near the conference room in which they will meet, which itself is a team room dedicated to the group not merely for the duration of their meeting, but the entire length of the current phase of their project. A project coordinator is the only permanent inhabitant of the area, responsible for various group tasks on-site. Upon reaching the entrance to the area, each participant pulls out their "rolling office"—in effect a small file tabouret on wheels containing personal files and personal effects—and tugs it into their assigned office. In the office is a desktop workstation and a docking station for the visitor's laptop PC.

The meeting room is similarly equipped with rich communications links in order to accommodate other participants who were not able to adjust their schedules to be physically present at the meeting. One wall is a "video wall", a large-size screen capable of being used for videoconferencing or displaying computer-mediated information. At the group's option, it can be used to display a remote colleague, display CAD information or display any of several feeds from remote sources.

Of course only a fraction of the Pentagon is now used by military personnel. The Secretary of the Air Force and her Secretariat remain at the Pentagon to conduct the business of the Air Force. Senior military leaders of the Air Force spend more time then perhaps they might like handling interactions with Congress, OSD, the Joint Staff, and occasional non-electronic meetings with the leadership of the other services.

15.0 Business Applications

15.1 Scenario

It is May 20, 2025. Major Drew Englund makes his way to the office. He is about 20 minutes ahead of his normal schedule. "I better let George know I am going to be early," he thinks. His personal communications device is on the seat next to him. He speaks. "Call George."

Instantly George replies, "Good morning, Drew."

"Good morning, George. I am going to be about 20 minutes early today, can you be ready for me?"

"The coffee will be hot and your mail ready for reading."

"Thanks, good bye."

While Drew completes his morning commute, George, his electronic assistant sets to work. George starts the automatic coffee machine and then begins to sort the morning's e-mail and generates an updated agenda for the day based upon two of the mail messages. As Drew arrives into the parking lot the Headquarters building comes alive, literally. It instantly recognizes Drew as an authorized occupant and opens the doors for him and turns lights on in the hallways as he makes his way to his office. The temperature in his office is carefully monitored by the building. Drew grabs a mug of coffee and sits at his desk.

"Good morning George, did you have any trouble getting everything together for me?"

"No, I didn't. I used your standard sort on the e-mail, are you ready for me to read them?"

"Yes, please go ahead."

George reads the five messages pertaining to the day while Drew reviews his calendar. "I blocked time for the two meetings and I prepared a draft presentation for the status review. Shall I confirm your attendance?"

"Yes and bring up the presentation. It's a good thing I am early today."

Drew reviews the draft presentation and adds his own touch to it. He then goes down the hall to the general's office. Most briefings are done by video but the boss requested his presence for this one. Drew wonders why. In the general's office he pulls up the briefing and talks the general through it. He concludes with a note about the test flight in the afternoon and the engineering review afterwards. "Our contractor is ahead of schedule and wants you to approve his performance bonus after today's test flight."

"Good, if it goes well today I will give you the authorization codes to make the transaction."

"Thank you, sir."

Drew heads off to lunch early to give himself time to prepare for the afternoon flight. He reviews the flight characteristics and the design specifications of the new dual-role fighter. He then finalizes the simulation scenario and flight plan for the test. The new fighter project was still a year away from building the first prototype but all indications pointed to being ahead of cost and time schedule. Drew headed to the VR-flight room for the test flight.

The VR-flight room was a new creation proving to be extremely valuable. The room was a melding of technologies enabling the pilot to “feel and fly” the aircraft under all conditions, peace or battle. He took his seat in the center of the room and the simulation began. The aircraft formed around him and the test scenario began. As he took off, he again experienced the thrill of flying that first attracted him to the Air Force. The test flight lasted almost two hours. Drew was tired but exhilarated by the capabilities of the new fighter. He attended the video debrief and engineering review to pass on his comments. Some tweaking needed to be done on the instrumentation displays and some major work needed to be done on the seating ergonomics. The designers concurred and promised new drawings would be sent by week’s end. The contractor manager asked about the bonus and Drew agreed they met the criteria and promised the transaction would be complete right after he received the next set of drawings.

15.2 Business Applications

The above scenario clearly depicts the level of technology growth we expect to see in the next 25 years. It highlights some of the technologies the Air Force can expect to use.

While most of the focus of the Information Technology Panel is on military operations, some consideration must be given to the business side of the Air Force. The business side is loosely defined as those daily operations not directly tied to prosecuting war, i.e., information management, legal, acquisition, personnel, finance, facilities management, environment, legislative liaison, public affairs, etc. These are the same functions found in most Fortune 500 companies, the overhead for doing business. These functions comprise much of the work the Air Force accomplishes on a daily basis during peacetime, contingency operations, and wartime situations. A substantial amount of commercial work in information technology focuses on these tasks, for obvious reasons; there are lots of paying customers with problems to solve.

The business side of the Air Force cannot be neglected in the New World Vistas study for two reasons. First, it comprises such a large part of Air Force peacetime - and wartime - daily functions. Second, industry is producing new products, new solutions, and new technologies every day that can provide immediate benefit to the Air Force and they are doing this at increasingly reduced cost. So, there is enormous cost savings potentials for the Air Force with the right information technology investment strategy.

15.3 Assumptions about the Air Force Structure and Business Practices

Before we consider the technologies that will impact the Air Force, we need to address some assumptions about how the Air Force will be structured in 2025 and the business practices that will be in place. The old adage “we will be joint, where joint makes sense” will be more true than ever. The business side is where joint really makes sense because all four services are structured like big business in many ways. We manage large payrolls, have vast amounts of real estate, perform property or facilities management on this real estate, manage health care plans, procure supplies and equipment and the list goes on. We will undergo a transition over the next several decades. First we will “go joint”. The services will combine their many similar business functions into joint systems using the same concept that is being used to create the Global Command, Control, and Communications System. These joint systems will rely more and more

on COTS products, both hardware and software. After this period, we will, for downsizing and cost savings, start contracting out these same services to commercial industry. Sometime around 2025 most of our business functions will be outsourced to the commercial sector. The first step towards this is already happening with American Express charge cards provided to government travelers instead of cash advances. There is no reason why most of our personnel accounting and finance functions could not be handled by a very competent vendor of outsourcing services.

Not all Air Force business functions will be privatized. There are several critical functions that are peculiarly military in nature or are mission critical functions, for example, logistics to the forward deployed forces. Lack of privatization does not mean these functions will be immune from commercialization. They will become extremely reliant on COTS technology. Sometimes the technology will need to be adapted to a military function but many times it won't.

It does not matter which functions are privatized because regardless of the function, some basic things take place in every office, Air Force or otherwise. Schedules and calendars need to be kept current. Phone calls happen. Documents - forms, letters, memos, and presentations - are created, proof-read, and transmitted, either by US Mail, electronic mail, fax, or VTC. Information is received from many sources like: TV, radio, magazines, newspaper, letters, electronic mail, or the Internet, to name just a few. There is not a noticeable difference in activity from office to office, whether Air Force or corporate. The data and purpose change from office to office but not the activities.

The Air Force is like any large enterprise and even though some functions have a military flavor, corporate America is doing something similar today. The Air Force functions need the same robustness, no more, no less, as private industry. We will not argue which functions will be contracted to private industry and which will remain blue-suit functions. We will assume a large number will be contracted.

Information technology will change the business world, therefore, it will change the Air Force. Corporate business will drive advances in information technology which in turn will change corporate business practices. The business side of the Air Force will follow these changes but will not drive the advances. Let's look at some of the key technologies and the impacts they will have on the Air Force in 2025. All of these technologies are discussed in greater depth in other portions of this Panel's report. These are just highlights impacting business practices.

15.4 Technologies Affecting Air Force Business Applications

Effectively infinite bandwidth, no more worrying about sending or receiving video or other bandwidth eating applications; this will happen through a combination of advances in multi-path fiber optics, ATM, and compression techniques. These three technologies are really only just starting to mature today. Effectively infinite bandwidth opens the door to many applications. Telecommuting, by 2025, will be telepresence, the ability to be where you really want to be and still operate in a very real way. Voice, video, data, or any combination will be available in real-time without distortion or sacrificing quality of one for the other. All this leads to the distributed workplace that will allow the Air Force to use smaller facilities saving money and improving morale. Video will be the communications means of choice for most applications. Conferences, meetings of all sizes will be via video as a matter of course without the distortion we tolerate today. Bandwidth on demand for any application at any location will do away with

the circuit-centric communications networks currently used. Commercial providers will meet almost all of the Air Force communication transmission needs because they can offer it more cheaply, more reliably, and just as secure as doing it with military resources.

There are two stepping stone technologies the Air Force will use over the next decade or more to get to a paperless environment. The first is the smart card. The current ID card is being replaced by a smart card carrying personnel, dental, and medical records. The card can also be programmed to contain security clearances and security codes for authentication and access. Instant access to vital information without the mess or bulk of paper.

Another stepping stone to a paperless Air Force is Electronic Commerce/Electronic Data Interchange (EC/EDI). EC/EDI will be ubiquitous in the 2025 marketplace, including the Air Force. Electronic transfer of contracting information and money will change and streamline the Air Force acquisition process.

Raw computing power will continue to grow and impact several technologies that will greatly affect daily operations. The interaction between man and machine will become more and more natural. Speech and gesture will replace keyboards and mice. Today's speech recognition engines will grow and mature allowing total control of the computer including the ability to dictate at a natural talking pace. Work production will increase because the built-in delays between thinking and typing will not exist. Typing a password will be replaced by the computer automatically recognizing authorized users using biometric methods. Offices where workers converse with their computers will be commonplace. No longer will operating a computer require a series of complex keyboard, mouse, and control character movements.

Computing power and advances in artificial intelligence will provide smart electronic assistants. These assistants will provide tremendous support to all levels, from the airman to the general (just like today's administrative assistants). The intelligent agent will screen electronic mail and phone calls. Other agents will do research, or find individuals and documents, even provide help in creating documents and briefings. Agents will help the action officer build briefings, including the difficult part of defining content and goals. Agents will be the "super-secretaries" of 2025. The agents will be personalized and take on personalities as they interact with their "masters".

Collaborative planning will fall out of very large bandwidth and increased computing power. The ability to share documents in real-time and communicate using voice and video simultaneously will be the norm. Action officers will be connected to others around the world just like they are sitting side by side. Operations centers can be distributed, with all the participants remotored from different geographic locations. IT-enhanced engineering, and data sharing allow weapons, buildings, planes, and other artifacts to be created better than ever by a distributed team of architects. Organizations can be smaller because expertise is only a "virtual access" away. Collaborative planning really integrates all the pieces making everything transparent to the user.

Smart buildings are an immature technology today that will grow significantly. This is an offshoot technology from the faster, smaller, better, more powerful computing technology. It merges high power computer chips, smart sensors, expert systems, and advanced processing to create a building that is an information appliance. The building can monitor and affect lighting, heating, cooling, and security appliances. Rooms and hallways are monitored for movement

and temperature and lighting is appropriately affected. Entryways are monitored and persons may be admitted or prevented from entering based upon a scan of physical characteristics like: iris, handprint, fingerprints, voice, and so on. This translates into significant savings in building maintenance costs and security personnel. It affects the business part of the Air Force because it affects costs and because good ergonomics makes employees comfortable and productive and saves health care time and money in the long run.

Modeling and simulation will have the biggest impact on the Air Force acquisition process. Integrating modeling, simulation, IT-enhanced engineering, and high powered computing will produce VR-rooms where virtual reality creations are life-sized "real" entities. These rooms will allow flight testing of aircraft, weapons, battle plans, room design, building design, and etc. Nearly perfect designs will go to producers because of the extensive testing, changing and additional testing that will occur in these rooms. It will be reconfigurable and reconstructable quickly. Intelligent agents will help the average Air Force action officer take advantage of these capabilities. Action officers will be able to test doctrine, battle plans as well as design plans using the advanced modeling capabilities that will exist.

Personal communication devices will provide an electronic tether wherever the action officer goes. They will provide voice, video, and data transmission capability, taking advantage of bandwidth on demand. These communication devices will allow the action officer to communicate with his electronic assistant wherever the assistant resides. Combined with speech recognition technology, this device will be simple and intuitive to use.

15.5 Summary

Corporate Air Force will be like the rest of corporate America. Information technology will affect daily business operations in many ways. Travel will be reduced. Geography will no longer limit planning capabilities. Acquisition will be based on modeling and simulation studies and IT-enhanced engineering. Business will be affected by EC/EDI enabling competition that can be accomplished in days not months or years. Natural interaction between man and machine will provide increased production and reduced frustration levels. Finally, Air Force personnel will be on electronic tethers with a multi-purpose personal communications device. Replace the uniforms with business suits and there will be no way to distinguish the Air Force "business" from other Fortune 500 corporations.

16.0 IT and Air Force Organization and Education

16.1 Introduction

The revolution in information and IT will profoundly affect the organization and education of the Air Force. As the industrial revolution created entirely new organizations of mass production and public education, the information revolution will lead to new forms of organization and education. Today's Internet provides an intriguing example of how unusual this new type of organization will appear. The embodiment of complexity, the Internet does not appear to be organized in any traditional hierarchical sense. But this system is organized very economically, and adaptively. A number of large organizations are beginning to learn about these new forms of organization, and the Air Force will inevitably, and appropriately, follow their lead. At this point these nascent ideas provide only azimuth, not destination. A final solution for a new Air Force organization can not now be known. What is known is that it is better for organizations to begin to become adaptive, rather than wait until more is known. As General Gordon Sullivan said,

“In this environment, the payoff will go to organizations which are versatile, flexible, and strategically agile, and to leaders who are bold, creative, innovative and inventive. Conversely, there is enormous risk in hesitation, undue precision, and a quest for certainty.”

Three general observations underlie our observations. First, IT affects all organizations, not just those in the computer industry. As air power once significantly changed army and navy doctrine and organization, IT changes every organization. It will change the Air Force even if the primary responsibility for the Information War mission is consigned to another agency. Second, IT will not win wars. Innovation of new technology and tactics give advantage only if well integrated to the organization's culture, policies, and doctrine. IT is just one system of an organization's “system of systems” to borrow the phrase from Admiral Owens. Third, organization and education will be vital differentiators for the Air Force in a copy-cat world of bits and bytes, but what are the characteristics of effective organization and education 30 years from now?

16.2 Organization

The beginnings of an answer can be found in organizations that are now successfully riding the IT bow wave. One of the most common observations about these enterprises is that they are relatively flat. Multiple levels of middle managers are not needed when IT is appropriately employed.

Prior to networked computing, one role of middle management was overseeing highly defined, discrete functions within a division and reporting on those findings. Responsibility for oversight had to be limited to specific domains as most information was not readily available. The information was known by the manager's people, down the hall on another desk, or in last year's folder. The IT revolution changes that. Middle management's span of oversight greatly expands as most of the important data is abstracted to information that can be quickly amassed and analyzed. Fewer middle managers are needed.

Another factor contributing to flatter organizations is the power of IT to support the coordinating, synchronizing function of managers. This too will change. Many scholars and practitioners see IT as a revolution in coordination not production or computation. Networked

computers are able substitutes for the coordination role as email, architecture, and other programming techniques dutifully convey top management vision to lower levels and relentlessly scan data for exceptions worthy of a coordinated response by management.

To effectively adapt the Air Force must move beyond its current monolithic, highly structured, chains of command. Designed to overcome attrition and the fog of war, redundant levels of command are increasingly inappropriate in an age of sensors that peer through the fog and create volumes of abstract data patterns ripe for accumulation and processing by high-level managers. Existing functional chains of command can be replaced by *networks* of command to reduce time and improve coordination of large numbers of forces. If we are to win at coordination, in a "Wal-Mart War" as General Marshall calls the next conflict, it will be due to using IT as Wal-Mart does to coordinate consumer demands and supplier production without delays.

The Air Force has already flattened by removing the Air Division level from most commands. The resistance to that change indicates how difficult continuing this process will be. A first step in a series of actions should be to develop an on-going industry education and touring program for flag officers to see how Wal-Mart and others are winning their IT wars. That group will produce an ongoing stream of necessary improvements. A second step will be to increase reward structures for organizations within the Air Force that demonstrate appropriate IT uses in this area.

Another quality of successful organizations is that they have learned how to produce hundreds of niche targeted products and services. IT permits adaptive, flexible manufacturing and delivery technology that permits organizations to offer hundreds of related products to more precisely satisfy increasingly specific demands from the marketplace for custom tailor-made products. Using database technology and consumer demand data, ever finer details of market niches are created for companies to pursue. As a result, companies like USAA offer hundreds of insurance and financial products to their specific customer base, and plan on delivering tailor-made products for every consumer in the future.

Analogously, the Air Force must adapt and learn how to strike and defend with hundreds of varied lethal and non-lethal weapons from physical ordinance to abstract bits and bytes. In a CNN communications world, hostile action with lethal weaponry is becoming unacceptable. The Air Force must continue to reorganize away from one major product—the delivery of lethal ordinance—to Global Presence, where presence is satisfied on a more "niche" basis. We must organize to deliver a vast array of coordinated products into the day-to-day, undeclared-war marketplace of today, as well as preserving the ability to strike an opposing force in the field.

This will be difficult in the near term for two reasons. The Air Force is currently dominated by operations, operations that have proven recently to be very effective. A second limitation is that we have no captivating vision of a post-industrial organization. First, the Air Force can undertake doctrinal review, determining what future "weapons" and methods of employment will be needed. Second, the Air Force should continue to expand the number of weapons and objectives used at Flag warfighting exercises at Nellis.

In addition to flat and niche driven, a third quality of successful organizations is that they use front line employees in a different way. These workers are empowered with intellectual and informational skills, and the power to act on what they see and observe, that is, what they learn.

As mentioned, IT generates ever finer details and tightens the link between top management vision and front-line employees. Both factors demand smarter, much more capable front-line employees with a primary task of noticing cues in the environment and communicating them. Linked to state-of-the-art IT, Frito-Lay trains college graduate route delivery salespeople to look for consumer purchasing patterns at specific outlets and authorizes them to act to adjust delivery and promotional plans. These well trained, autonomously acting front-line employees input this ever changing data into an IT system so that patterns in the small changes can be noticed by others and production and marketing plans of specific products reprogrammed. "Constantly learning" organizations employ highly educated front-line workers and attempt to leverage their skills to out-adapt their competition. These efficient adaptations to subtle market fluctuations replace the old notion of organization as machine that scales mass production up and down.

To do this ever finer tuning, organizations need smarter front-line workers. How the organization accomplishes this education will be discussed later, but the goal is a front-line work force with new skills. New intellectual skills will enable workers to create meaning and value from the abstract cues of data and information. Abstract thinking, problem solving, inference, explicit understanding of work, modes of reasoning that are analytic, procedural, and the ability to commit attention to mental effort: these skills permit front line workers to separate truth from disinformation, recognize actual threats, and understand how their roles must adapt.

Currently, the Air Force is decentralizing the power to act on decisions, but does little *intellectual* and *informational* skill education necessary for sophisticated learning. Instead it tends to rely on training that emphasizes standardization, top down direction, on-again off-again phasing, and physical skill training.

Again, the multi-level hierarchical model must change. A better model might be a network. Smart front-line employees at the periphery, managers as routers and switches and strategic planners as network designers. The main source of this network's adaptiveness and resilience is the smart decentralized "end of the network".

16.3 Education for the IT Era

A first step might be to alter educational objectives for undergraduate and first level training programs for our front-line members. This leads naturally to the topic of future Air Force education.

New educational processes will become commonplace in the next 30 years to satisfy the needs of these post-hierarchical organizations. From primary to higher education, the educational processes that provide our young airmen and officers will make far greater use of IT. As a result, our newest members will have increasing capability and fluency with IT skills. Once on active duty their education will become continuous, and applied. Education will grow in scope beyond its training aspect to be more broadly defined as intellectual and information skills such as pattern recognition and inference. Current skill training programs will continue, but the emphasis will shift to applied education that will become a vital daily transaction for the organization, on equal footing to sortie generation.

Continuing education will be explicitly fitted to workplace needs, a tailor-made service that the academic people will learn how to do. Work and learning will become indistinct. Learning will replace labor as the only sustainable source of advantage.

For the Air Force, education and training may be examined in three levels. For front-line workforce, one idea could be contracted mentor-nets involving hundreds of the best educators and thinkers in industry and universities they will provide the opportunity for Air Force "cyber-pilots" to hone skills using the latest technology, share insights, notice threats, exploit adversary weaknesses, and differentiate our capabilities from hostile organizations. Student tribes may be constantly engaged in mock cyber-fights simulating deployment of aircraft, hostile acts, intention masking, and disinformation passing.

For managers tasked to educate and integrate, this "education net" aspect of work would grow to include education and teamwork concepts. To extend the example, managers might be given opportunity to purchase untried technology for their tribe. Those managers that can quickly organize their unit over a network to weigh alternatives and applications might be rewarded for speed and integration of technology.

Top leaders would continue their education at a reduced pace. Once more the previous example might entail these top leaders selecting scenarios and rule structures that they feel appropriate to shifting national security climate and need. The goal of this process would not be to command, but to mold organizations to become adaptive to front-line input.

The old adage rings true. You fight the way you practice. We must be engaged in continual learning in the space in which we will need to attack and defend.

16.4 Conclusion

Organization and education are not distinct, but blended facets of information power, a prerequisite of successful organization. Reinforcing each other they help create a combined system capable of fighting and defending in hundreds of niches. This is similar to how smart terminals within simple networks have replaced large central controlling computer networks. We must become a mesh of brilliant "end points" of an information network.

17.0 Epilogue: Principles to Think With

Epilogue

Woven through this long report on the future of Information Technology, hovering above the details, have been some principles with which to think about the future of IT. The future may differ in detail from what we predicted, but the principles will not. We found a statement of these “principles to think with” written out in the 1993 book, *The Road to 2015*, by the futurist John Peterson of the Arlington Institute. With his permission, we reproduce them below.

- **Speed.** Everything is going faster, so speed is increasingly being used to measure value.
- **Trend toward light.** Light is the fastest communications medium with the greatest capacity to carry information. There is a clear trend toward using light and optics in information technology.
- **Information.** Information, in the form of knowledge, is what allows speed. Information is the capital commodity of the future.
- **Going to digital.** Information in all forms is being converted into digital forms so that it all can be transmitted through the same cables, fibers, frequencies, and equipment.
- **Global connectivity.** Everything is being connected to everything else. In time, almost every home, office, school, and government agency in the developed world will be connected in a huge information system by way of every computer and telephone.
- **Global accessibility.** There will be no place on the surface of the earth where one can't access the whole network.
- **Moving information instead of people.** Information technology is making it more advantageous—in almost every situation—to move information rather than people.
- **Power migrating toward individuals.** Individual people will increasingly be able to access, analyze, and manipulate information (the source of wealth and power) without the need for larger organizations like corporations and governments.
- **Systems thinking.** All things of importance are coming to be understood as systems; in most cases, they are highly complex, dynamic sets of sometimes widely dispersed components. In science, particularly, there is a move toward the integration of disciplines. In manufacturing, there is a move toward concurrent engineering.
- **Increasing complexity.** Manufactured systems are becoming more complex and faster.
- **Increasing vulnerability.** The more complex a system becomes, the more likely the chance of system failure. There are unknown secondary effects and particularly vulnerable nodes.

- Qualitative becoming more important than quantitative. Software, intuition, speed, and quality are areas that are pregnant with opportunity.
- New structures and organizations. All structures and organizations (business, government education) will reconfigure to adapt to the faster, more interconnected world and to the more powerful and enabled individual.
- New info-criminals. There will be a significant increase in information crime, more viruses, and a growing international information criminal element.
- Unpredictability. Technology is a huge effort with lots of people in an extremely complex context where there are fundamental changes in the underlying principles. We almost certainly will not anticipate some of the significant changes that will certainly occur.
- Punctuated change. One or more of the reality-exploding nascent technologies will come to fruition soon and send shock waves throughout the global system.

Appendix A

Information Technology Panel Charter

Information is an abundant resource collected and disseminated many times without direction. Information is the controlling factor for the future. Information technology allows mankind to control, direct, and use the vast amounts of information available. Information technology is undergoing an explosion of growth that will continue into the 21st century. As part of the *New World Vistas* study for the Air Force Scientific Advisory Board the Information Technology Panel will focus on the impact of information technology on tomorrow's Air Force. How does the advance of hardware, software, communications, and all the associated disciplines, and applications produce advances in the military operations of the Air Force?

Appendix B

Panel Members and Affiliations

SAB Members

Dr. Edward A. Feigenbaum, Chairman
Chief Scientist, USAF

Dr. Barry W. Boehm
Dir., USC Center for Software Engineering
University of Southern California

Dr. Randall Davis
Assoc. Dir., Artificial Intelligence Laboratory
Massachusetts Institute of Technology

Dr. Robert W. Lucky
V.P., Applied Research
Bell Communications Research

Dr. Donald L. Nielson
Dir., Computing and Engineering Sciences
Division
SRI International

Executive Officers

Major M. Clarke Englund
Chief, Technology Insertion
HQ USAF/SCTT

Capt Dean Osgood
Executive Officer
USAF Scientific Advisory Board

Research Assistant

Lt Tracy Adams
Graduate Student
MIT

Participants

Mr. Paul Saffo
Institute for the Future

Col Roderick Taylor
Chief, Technology Division
HQ USAF/SCTT

Dr. Gio Wiederhold
Computer Science Department
Stanford University

Major Earl H. McKinney
HQ USAFA

Appendix C

Forecasting Tables

These tables are the result of brainstorming by the Panel members during the summer study. The goal of the session was to identify the most significant happenings in each of the focus areas and explain why they were significant to the Air Force. Several focus areas did not get discussed due to time constraints.

Communications

What?	When?	Why Important?
Most important applications for comm are surprises to authors of this report and to everyone else.		History shows that, although we can predict infrastructure, <i>we cannot predict applications</i> ; don't use today's applications to plan future needs for capacity.
Data traffic on all nets continues to double annually.	1995-2010	Fundamental law of nature; enables projection of capacity requirements in spite of unknown applications.
Global commercial optical networks, 10Gbps per wavelength, per fiber, wavelength division multiplexed.	2000	Ultra cheap ultra-high-speed backbone gives essentially infinite bandwidth.
Broadband commercial satellite network (>1Mb) (e.g., Teledesic)	2005 (.5 probability)	Commercial broadband comm in remote regions.
At least one global commercial satellite telephone network exists (e.g., Iridium).	2000 (.8 probability)	Limited capacity voice comm available everywhere.
Widely available commercial wireless infrastructure in developing nations.	2005	Usable commercial infrastructure almost everywhere for voice and low data rates.
ATM becomes standardized broadband environment.	2000	COTS availability of broadband computer comm equipment.

What?	What?	Why Important?
IP v6 (or vx) is standard computer protocol.	2000	Commercial basis for real time signals (e.g., video conferencing) on packet nets, policy-based routing, security, and expanded address space.
<i>Internet as global information infrastructure .5B people connected using cheap, ubiquitous "Internet phone/data appliance."</i>	2000-2005	<i>Overarching, super-national civilian connectivity profound social, business and military consequences; intel consequences.</i>
Computer network warfare	2000 (depends on investment)	Delicate balance between network attacks and defense; critical differentiator
World Wide Web becomes collaborative, with virtual reality, interactive, point-of-view browsers.	2000	This will enable the AF to find new ways of conducting its business and operations.
Number of large companies decreases and virtual companies multiply because of networks.	2000	AF may have to remodel itself organizationally in radical fashion.
Commercial comm facilities become cheaper by factor of 10 (Over and above normal price decline of 11% annual) because of adoption of Internet paradigm with complexity and cost at periphery.	2005 (.4 probability)	Makes AF comm facilities that don't rely on commercial facilities uneconomical.
Instant broadband wireless infrastructure; wireless LANs, deployable distributed wireless self-organizing nets.	2000 (depends on investment)	Deployable broadband infrastructure.
Unattended autonomous vehicles used to fly patterns for comm relay.	2005	Good solution to the 100-mile gap between forward units and commercial gateways.
Broadband comm links into tactical aircraft provided by satellite.	2005	Bandwidth for sensor telemetry and commands to the smart or remotely piloted vehicle.

What?	When?	Why Important?
World standards for encryption with wide commercial usage.	2000-2005	Unbreakable codes for both sides; military can use commercial security packages.
Agents negotiate QOS (Quality of Service) across heterogeneous nets.	2000	Buy just what you need for a given application (e.g., video) from multiple vendors.
There is a breakthrough by a factor of 10 in compression efficiency.	2000-2005	Demands for bandwidth will eat up the gain.
Multimedia queries can be accessed by content.	2000-2005	Find pictures and sounds by description of content.
Global knowledge banks codify "knowledge" in many fields.	2000-2005	Better than libraries or databases.
Electronic agents, filters, and communications surrogates mediate communications.	2000	More efficient communications—electronic "secretaries."

Human-Computer Interaction

Computer uses are so broad and definitions and perceptions of functionality so user-dependent, that development metrics such as degree-of-intelligence are very difficult to define. So, the development of human-computer interaction should be thought of principally as a direction and not an end. (Dates are 2-4 years after introduction to commercial marketplace.)

What?	When?	Why Important?
Intellectual Support		
Natural language understanding - (Text, speech, writing and gesture recognition plus understanding but separately employed.)	2000	<ul style="list-style-type: none"> - Ultra-mobility - Hands-occupied operations - Task-specific needs
Modality integration - (All interaction components taken in combination appropriate to situation.)	2002	<ul style="list-style-type: none"> - Task adaptation - User preference adaptation
Task-oriented commands - (Interaction above present application level. Context-aware tasking and execution.)	2005	<ul style="list-style-type: none"> - Work-centered, context-aware interaction combining the offerings of present-day applications programs - Visual programming at task levels
Delegation - (User-like understanding and reaction to tasking, dialogue)	2010 (w/investment)	<ul style="list-style-type: none"> - Task machine in natural language to carry out goals - Intuitive, training-free interactions; one-session results
Hardware: - Displays: -- Flat, low power, print-quality (300 dpi ATLC) -- Wall (projected) with arbitrary resolution -- Bright, flat, high resolution, low power	1998	<ul style="list-style-type: none"> - For desktop and mobile applications - Projected wall-size seamlessly tiled displays with very high, "walk-up-to" resolution

What?	When?	Why Important?
<ul style="list-style-type: none"> - Scanners - smart paper (encoded background) - Speech - Continuous, speaker-independent, 1-2% error rate, correctable with context - Handwriting - (same as speech) - Video conferencing - (See Communications.) 	<p>1998</p> <p>2004 (w/investment)</p> <p>1997</p> <p>2000</p>	<ul style="list-style-type: none"> - Tells scanners details of digital page conventions. - Speech and handwriting modules available and affordable in server, workstations, PCs, etc. User toolkits.
<p>Virtual Reality</p> <p>Totally synthesized environments that are sensory-matched to human capabilities. Consists of head-mounted stereo visual and aural presentation, a tracking system, a controller, and a haptic output.</p>	2005	<ul style="list-style-type: none"> - Visualization of other, altered worlds - Training scenarios/simulation [Caution: until sensory mismatches are corrected, VR systems may cause sickness under long term use]
<p>Hardware</p> <ul style="list-style-type: none"> - Displays - Need head mounted, 3D, very high resolution ($\sim 10^{10}$ pixels), object-specific focusing (Display not fixed focused) - Trackers - Need 6D wireless, low hysteresis - Controllers - Need speech recognition, gestures 	2005	<ul style="list-style-type: none"> - Entertainment market may deliver affordable versions.
<p>Augmented Reality</p> <p>Overlaying synthetic cues and structure over real-time, real-world activity. May require transparent, but imaging glasses and some tracking system with spatial synchronizing of real and synthetic images or a projected spatially synchronized overlay on the target object.</p>	2005	<ul style="list-style-type: none"> - Training, performance evaluation - Maintenance tasks - Hazardous undertakings - Simulations - Constraints/enhancements can alter real capabilities

What?**Hardware**

- Requires transparent glasses which are also displays.
- Trackers - Need to be an order of magnitude more accurate than VR trackers, depending upon distance separating real and overlaid image.

When?

2005
(w/investment)

Why Important?

- Rapid advance of inductees' skills alone would be very profitable in above roles. Fewer airmen to carry out increasingly technical roles. This technology may not get developed by the entertainment industry.

Telepresence

Real-time/human sensing and physical dexterity are translated into inaccessible spaces.

2005
(w/invest)
(evolutionary)

- Remote operations (some delay-bounded): RPVs, surgery, hazardous ordnance, viewing and sensing,....
- Motion translation: minimally invasive surgery,
- Scale transformations: micro-system assembly/repair, satellite and space-station repair
- Distributed conferencing with extremely high realism

Information Access Technology

Definition: Data to Information flow support.

Situation awareness and decision support:

Post sensor, public, intel data acquisition. Pre display and HCI.

Supports smart processing technologies as Planning, S&M, ...

What?	When?	Why Important?
Information integration: Support mission planning Support analysis of multiple future courses of action Dynamic integration of actual events, replacing projections with facts	2000 2010 2015	Mission execution, rapid replanning
Image indexing technology: obtaining and selecting relevant intel information (spatial, temporal selection) *semantic contribution to indexing.	2005 (some COTS, costly technology)	Volume of image-based input growing. Helps <i>avoid information overload</i> .
Data reduction by intelligent selection and compression (delta transmission) for limited bandwidth links.	2005	Reduce bandwidth for the last 100 miles, build on COTS model based compression research.
Composable software to process information	2000 (within one standard), evolving	COTS and standards acceptance dependent.
Architecture for rapid composition of information systems with broad access capability.	2015 (improvable by investment)	COTS, but focused on uni-directional interoperation
Tools	2005	Verification (OT&E) of information flow for reliability, trust, etc.
Mediators: a form of service agents (SW nodes, people for maintenance) to process potentially relevant information to create value for customers.	2000 (and evolving)	

What?	When?	Why Important?
Security mediators - owned by security officers, serve as intelligent gateways.	2005 (and evolving)	Needed to interoperate among domains at differing security levels.
Facilitators: locate potentially relevant information, from public and commercial (non-secure) sources.	2005 (Internet) 2015 (arbitrary nets)	Needed to exploit all available resources, some DoD adaptation.

Intelligent Software Agents

Intelligent agents:

- Are subordinate to a user or owner.
- Have degrees of autonomy.
- Are goal and not means driven.
- May replicate themselves.
- May have only qualified roles in safety-critical situations.
(Dates are 2-4 years after introduction to the commercial marketplace.)

What?	When?	Why Important?
Agents by Type Directed-action agents ("Do-this" agent) - Has fixed goals but can react to the data and data sources it encounters. Example: Uses the relational WWW search sites and routines to retrieve stated objects (subjects).	1997 (evolutionary)	<ul style="list-style-type: none">- Retrieve information in a well-indexed world- Sorting information flow- HCI assistance such as speech or network access- Framework for modularized intelligent systems- Collaboration with fixed goals and guidelines
Reasoned-action agents ("Achieve-this" agent) - Has fixed goals and is able to monitor other objects (data and processes). It can also reason about what it "sees" and take appropriate action. Continuing with above example: also examines the relationally derived sources and evaluates source attributes according to general guidelines. May prioritize sources given subject and recall for future use. Negotiates under strict rules.	2000-2005 (evolutionary)	<ul style="list-style-type: none">- Advisory agents for evaluation and training- Planning agent- Personal assistant that reasons; i.e., an HCI facilitator for a group of directed-action agents- Information probes- Gives system resilience or fault recovery- Heterogeneous database mediation- Agent collaboration with task-sharing

What?

Learned-action agents (“Accept-this” agent)

- Has above capabilities plus accepts more general goals and is capable of altering or adding to them under guidelines. Continuing with above example: accepts general direction for information assistance; infers recent user interest in given subject; sets goals for possible information retrieval in that subject; sets up indexing to relevant WWW sources as well as locally created citations. All this to anticipate user need.

Agents by Functionality

Advisory agent - Able to evaluate human or system performance in a circumscribed, application-specific framework and make corresponding recommendations and perform goal-seeking actions

Personal assistant - Will handle many HCI tasks like email, conferencing, interaction modalities, information filtering, scheduling, etc.

Offers a broader representation of the user in permitted jurisdictions.

When?

2005-2015
(evolutionary)

1997-2005
(evolutionary)

2000-2005

2000-2005
(evolutionary)

Why Important?

- Very abstract tasking
- Personal assistant that has broad awareness from both operations and its own probing
- May use guidelines and inference to set new goals
- Agent collaboration with dynamic task sharing
- Context-sensitive help in HCI
- May facilitate or manage a set of direct action or reasoning agents

- Teaching and training aids
- Planning systems - offers tomorrow’s trial ATO
- Monitors maintenance records on jet engines to predict need for major overhaul or replacement
- Manages parts arrival in just-in-time scheduling

- Handles many well formatted and unambiguous tasks
- Finds and reserves next available flight

- Calendarizes some user activities
- Monitor and approve movement of consignments

What?

Traveling agent - Can transparently retrieve well-defined information from a variety of distributed sources [Relies on consistent library interfaces]

Eventually learns retrieval environment itself for subjects of stated interest. Decides whether to use remote or user's cycles (host) in pursuit of goal.

Will provide broad and transparent representation for all well-defined user transactions, filtering information streams for user interests.

Collaborating agents- Multiple, function-specific agents form basis for intelligent systems. The simplest act at the behest of a facilitator or coordination agent, the more intelligent work out task sharing, and the most intelligent do dynamic task sharing to meet goals.

Design methodology for intelligent system design, including how intelligence is distributed in a system and the most appropriate abstraction for component interaction.

When?

2000-2005
(evolutionary)

2005-2010
(evolutionary)

2010-2015
(evolutionary)

2005-2010
(evolutionary)

2010
(evolutionary)

Why Important?

- Retrieves well-indexed information given indices or well-defined subjects [Caution: As much as 90% of all stored data is in unstructured form and perhaps irretrievable.]

- Information awareness and retrieval with user-defined guidelines
- Heterogeneous database mediation

- Formulates and remembers preferences
- Avoid safety critical situations

- Modularization of computer-based systems
- Deliver integrated functionality for HCI
- In a distributed computing world, monitors operations to "learn" how crashing occurs, what processes are affected, and how to recover
- Basis for pilot's virtual associates

- A framework for intelligent, distributed system design
- Used in graceful degradation or fault tolerance where multiple functions exist

Artificial Intelligence

What?	When?	Why Important?
Natural Language Understanding: 95% in specialty areas 95% in areas of common knowledge and culture 95% across breadth of human experience	2000 2010 2015	For ease of use and broad applicability of computers.
Knowledge Web of tens of millions to hundreds of millions of pieces of knowledge, distributed over Internet.	2005-2015	Enabler for Intelligent Applications e.g., critical mass for modular learning e.g., intelligent systems with common sense.
Inflection point in knowledge acquisition, learning, adaptivity by logical and heuristic methods.	2005-2015	Because of its relation to domain-specific software tools, predict order of magnitude increase in speed of software development, for example for training, intelligent agents, and data fusion.
Synthesis of neural nets and symbolic reasoners.	Evolving	Use neural nets for pattern recognition and symbolic reasoners for understanding.
Learning based on evolutionary models.	Slowly evolving	No miracles.
Plausible, "soft", inexact, fuzzy, probabilistic reasoning methods will dominate traditional formal logic and perhaps all other forms of computation.	2015-2220	Most problems that computers will be assisting with are not exact, numerical, or Boolean (black-or-white). Example uses: battle mgmt. and planning, sensor fusion, decision support, risk mgmt.
Signal + Symbol information fusion	Depends on effort but needs ACTD	Situational assessment.
Design rationale management	Evolving	Makes redesign and maintenance both feasible and affordable.

Planning and Scheduling

What?	When?	Why Important?
Basic plan rationale capture.	2000	Semi-automated case-based trans. and force planning. Semi-automated plan modification One COA.
Advanced plan rationale capture.	2005	90% automated plan selection and modification, based on planner's goals; multiple COAs.
Decision theoretic control of inference.	2005	Adaptive planning in uncertain environments.
Modeling and simulation advances.	2010	Campaign planning via detailed simulation.
High speed constraint-based scheduling.	2000	Order of magnitude increase in speed of scheduling from TPPFDs.
Rationale-based schedule repair.	2005	Fast, high-level modification of schedules.
Coordinated multi-agent planning.	2000	Automated planning that takes into account multiple perspectives.
Concurrent multi-agent planning.	2005	Order of magnitude faster multiple perspective planning.
Continuous, mixed initiative planning teams.	2010	Planning as a continuous around the clock process, constantly updated by a team of people and software agents working together.

Modeling and Simulation (M&S)

What?	When?	Why Important?
<p>Modeling and simulation interoperability</p> <ul style="list-style-type: none"> - Basic: data definitions, declarative protocols, 2 levels. - Advanced: assumption reconciliation, multilevel dynamics, advanced media. 	<p>Basic- 2005 Advanced- 2015</p>	<p>Ability to capitalize on multiple M&S's for combat systems definition and acquisition, test, production and logistics, education and training. Requires continuing evolution of interoperability standards and conventions, application to M&S.</p>
<p>User languages for M&S generation and analysis.</p> <ul style="list-style-type: none"> - Basic: Assembly with simple interoperability support. - Advanced: Automatic M&S configuration based to satisfy user goals. 	<p>Basic- 2005 Advanced- 2015</p>	<p>Rapid, flexible composition of complex M&S.</p>
<p>Believable semiautomated forces (SAFOR) for full combat spectrum.</p> <ul style="list-style-type: none"> - Basic: reactive to current state, based on simple heuristics. - Advanced: complex heuristics; learning capabilities. 	<p>Demo quality- 1995 Basic- 2005 Advanced- 2015</p>	<p>Cost-effective M&S salability. Enable experimentation with doctrine and strategies.</p>
<p>Virtual system acquisition: flexible migration from virtual to actual system capabilities.</p>	<p>2015 - for mature domains (e.g. mainstream aircraft, missiles, spacecraft). 2025 - for complex domains (e.g. integrated combat systems including C³I, battle management).</p>	<p>Rapid, flexible acquisition of pre-exercised systems. Complementary closed-loop exercise and evolution of system and its M&S support. Requires strong M&S support, revised acquisition procedures, virtual competition ground rule, education.</p>

What?

Defense simulation Internet: Network support of "Louisiana Maneuvers" scale activities.

M&S verification, validation, and accreditation.

- Basic: Test suites, simple assertion checking.
- Intermediate: Simple domain-model and built-in-test checking.
- Advanced: Agent-based domain model checking and dynamic built-in test.

When?

2000
- occasional-
exercise support
2010- regular
use

Basic- 1995
Intermediate-
2005

Advanced- 2015

Why Important?

Network infrastructure for theater-level simulations, virtual reality.

Critical to M&S credibility. Requires significant investment concurrent with M&S development and upgrade.

Software Development Technology

What?	When?	Why Important?
Powerful User Programming. Over 50% of workforce able to harness packages with millions of lines of code in a few hours to accomplish a semi-custom design.	Corporate information management: 2005 Combat systems: 2005-2015 (depending on investments in domain engineering)	Users get the software they want right away (for straightforward applications). Requires education, safety control limits on user programs, another level of application interoperability conventions to be able to bring user program together.
Concurrent engineering of devices and their software through integration of CAD, CAM and CASE.	2005-2015	Better allocation of functions to hardware, software, people. Moving software development to earlier in the product cycle. Requires bringing together physical models and computing models. Treating software as a first-class system design element.
Adaptive software: systems that automatically improve themselves based on observation of usage and data patterns.	2005: used by > 50% of Fortune 100 companies	Rapid, cheaper, more assured software adaptation to changing situations (e.g., Serbian mobilization patterns). Requires countermeasure defenses.
COTS software componentry and CASE tools enhanced to meet Air Force combat needs: security, survivability, real-time performance, scalability.	Continuing	Air Force combat software needs satisfied by commercial capabilities. Requires pro-active Air Force and DOD COTS software evaluation, enhancement R&D, and stimulation of commercial enhancements.
Commercial penetration of ADA will continue to be low, due to C/C++ COTS critical mass.	Continuing	Requires ADA/COTS bindings and interoperability packages.
"Beyond-Object-oriented" programming language better than C++ will be developed.	2005	Opportunity to support more disciplined and flexible software engineering.

High Assurance Systems

What?	When?	Why Important?
Highly penetrable - Military Systems - Civilian Infrastructure	NOW	We are vulnerable, but basic solutions exist.
Cryptography routinely embedded transparently.	2000	AF Key policies should be developed. Key escrow should be used because cryptography will be unbreakable, attack points will shift (e.g., Human/Protocol).
Cryptography uses expand.	Slow Evolution	Most effort should be in integration of use.
Theory Molecular Computation Quantum Computation	KEO	If P=NP, new encryption methods will be needed. New models of computation may arise.
Embedded Multi-Source Biometric Identification.	2000	Basic security embedded in terminals, etc.
Unobtrusive Ubiquitous Continuous biometric identity.	2010	"Smart locks" widely available.
Basic Survivability Via Dispersal.	Now	One contribution to high assurance
Graceful degradation.	Depends on investment	Load balance between man and machine. Redundancy between man and machine (Optimization and degradation are opposed).
Rules of Engagement.		There are none, policy needed!

Organization and Education

Organization is a differentiator.
Conflicts are organizational phenomena.
Developments are non-discrete.
IT will change AF independent of IW mission (e.g., airpower).
Severe organizational culture limitations (e.g., IBM mainframe).

What?	When?	Why Important?
Paradigmatic organizational change e.g., flat architecture, beyond virtual reality.	2015	Assessment, control, etc.
Emphasis shift away from top down to become adaptive.	2005	Education, personnel systems, IW mission.
Profound educational change.	2010	Selection, continuing education/ training.

Open questions:

On-going process suggestions.
Enablers or other structural techniques.
Is the Air Force a reasonable organization to compete in this mission.

Appendix D

Evolution and Breakthrough

Table V offers another view of the IT future that the panel constructed. It is a view of the 10-30 year future composed of:

1. "Evolutionary successes" with high probability of being achieved
2. "Breakthrough possibilities" with low probability but revolutionary impact
3. "Issues" as technological, educational, and organizational concerns for the future of the Air Force in the era of the information revolution.

Through a complex process of proposing and voting, the panel created a list of 55 candidates, and narrowed it down to the list of 19 that are shown in Table V. Though the panel used this list to stimulate its own thinking, it saw this list as possibly of use to others.

Table V Evolutionary Successes; Breakthrough Possibilities; Issues

Evolutionary Successes:

- Smart materials: Materials will change shape, color, texture, etc., to fit different situations.
- Human-Computer Interaction will be based on interpretations of multi-modal signals. We will interface with our computers through speech, hand gestures, eye movement, etc., imitating some of the ways in which we communicate with other humans.
- Telepresence: What we do as humans will be transferred to other things, bigger or smaller. Telepresence acts for humans under human control remotely.
- Virtual pilot, virtual crew: One or more of an aircraft crew will be replaced by software on-board or by crew members operating remotely.
- Widespread use of COTS (Commercial-Off-The-Shelf): Air Force will outsource functions, use commercial services and products, integrating them with Air Force systems (versus internally building or providing functions).
- Data broadcast satellite: Like broadcast TV, large amounts of data will be broadcast to a wide audience.
- Continuous speech understanding: Dictation transcription, query interpretation, and other natural language understanding tasks will be handled with broad scope (not limited domains as at present).
- Logistics/acquisition: Computer-enabled "smart" equipment embedded in equipment will automatically track status, do diagnostics, maintain maintenance records, and assist maintenance.

- Data fusion: Many different sources of information, some being sensor data, some being symbolic knowledge, will be fused together into accurate situation models, assisting situation awareness.
- MEMS, micro-electromechanical systems: Electromechanical systems on inexpensive microships will perform a wide variety of sensor and effector functions.

Breakthrough Possibilities:

- High density memory: We will increase our ability to store large amounts of data in small volume (very high “bit density”) by many orders of magnitude using optical and holographic methods and perhaps by storing directly into molecules.
- I/O Processing: Most of “computing” will be the handling of inputs and outputs, e.g., video and sensor data.
- Thoughts to electric patterns: Sufficient correlation will be made between patterns of thinking and intended human action such that an input channel from human mental activity to computer can be fashioned (bypassing muscle action, speech, etc.).

Issues:

- Organization designed for information exploitation: Information will be an organization’s most critical resource, after personnel.
- Graceful degradation: Technology that fails will automatically reconfigures itself from redundant functionality (that is, as humans did on the Apollo 13 mission).
- Info-centric aircraft: Aircraft will be designed to maximize information processing potential (in contrast to 20th Century aircraft design that optimized aerodynamic, not information, performance)
- Architecture: Information systems will have a clear architecture in terms of design rules, interface standards, etc. and there will be families of systems with the same architecture (saving money and system-building time).
- Infowar: Who is the enemy? What are his capabilities? Is the threat economic, terror, war? Who defends/monitors cyberspace?
- Security in openness—the so-called “spotlight weapon”: Can we use our intelligence information in a more open way to carry out a “CNN-type” mission in which the pressure of nations and populaces are used to bring about a desired end (in lieu of either diplomacy or projected force)?